

GILA SHER

EPISTEMIC FRICTION

*An Essay
on Knowledge,
Truth, and
Logic*

OXFORD

Epistemic Friction

Epistemic Friction

*An Essay on Knowledge, Truth,
and Logic*

Gila Sher

OXFORD
UNIVERSITY PRESS

OXFORD
UNIVERSITY PRESS

Great Clarendon Street, Oxford, OX2 6DP,
United Kingdom

Oxford University Press is a department of the University of Oxford.
It furthers the University's objective of excellence in research, scholarship,
and education by publishing worldwide. Oxford is a registered trade mark of
Oxford University Press in the UK and in certain other countries

© Gila Sher 2016

The moral rights of the author have been asserted

First Edition published in 2016

Impression: 1

All rights reserved. No part of this publication may be reproduced, stored in
a retrieval system, or transmitted, in any form or by any means, without the
prior permission in writing of Oxford University Press, or as expressly permitted
by law, by licence or under terms agreed with the appropriate reprographics
rights organization. Enquiries concerning reproduction outside the scope of the
above should be sent to the Rights Department, Oxford University Press, at the
address above

You must not circulate this work in any other form
and you must impose this same condition on any acquirer

Published in the United States of America by Oxford University Press
198 Madison Avenue, New York, NY 10016, United States of America

British Library Cataloguing in Publication Data
Data available

Library of Congress Control Number: 2016938685

ISBN 978-0-19-876868-5

Printed in Great Britain by
Clays Ltd, St Ives plc

Links to third party websites are provided by Oxford in good faith and
for information only. Oxford disclaims any responsibility for the materials
contained in any third party website referenced in this work.

For Itai and Shlomi

*The light dove, cleaving the air in her free flight,
and feeling its resistance, might imagine that
her flight would be still easier in empty space.*

Immanuel Kant

Preface

In this essay I develop an integrated theory of knowledge, truth, and logic. I do this by tying together, and further expanding, three projects I have been working on since completing my dissertation in 1989: a model of knowledge that combines Quinean and anti-Quinean elements, a substantivist theory of truth, and a philosophical foundation for logic. What led me to embark on this extended project was the realization that its three parts were so thoroughly interrelated that I could no longer either adequately explain or continue developing any one of them without connecting it to the others. My model of knowledge requires a non-traditional correspondence account of truth, as well as a new foundation for logic. My foundational account of logic centers on issues of veridicality (truth), and the methodology that makes a veridical foundation for logic possible is a non-traditional methodology that was developed for the model as a whole. Similarly, my substantivist theory of truth is motivated by epistemic considerations associated with the model. It exemplifies the model's approach to philosophical theorizing, and it develops a non-traditional conception of correspondence that is flexible enough to encompass all fields of knowledge, including mathematics and logic. A common thread passing through all three projects is the centrality of both freedom and friction (constraint) to knowledge, including philosophical and logical knowledge. In this essay I put more emphasis on friction; in a sequel, *Epistemic Freedom*, I hope to further pursue the second aspect, freedom.

It is impossible to introduce all the ideas developed in this essay in a brief preface. Instead, I will offer a bird's eye view of the essay, focusing on some of its ideas and leaving others to the work itself.

Underlying my approach to knowledge is an interest in "the basic human epistemic situation". Among the central elements of this situation are: (i) We live in a world of which we are a part. (ii) We desire to know and understand this world not just practically but also theoretically. But (iii) our cognitive resources are considerably limited and this renders the world highly complex relative to our epistemic abilities. Nevertheless, (iv) we humans are ambitious creatures and we aspire to know the world in its full complexity. What makes our aspirations achievable (to some significant degree) is the circumstance that (v) while we are cognitively limited in some ways, we are cognitively endowed in others. Our cognitive capacities are intricate, they involve not just sensory perception but also intellect and imagination, they include the ability to participate actively in the project of knowledge: plan, initiate, design, shift perspectives, critically evaluate our epistemic strategies and results, create technological tools, and so on. In light of this complex situation, a central question of epistemology is whether and to what extent we are capable of

acquiring theoretical knowledge of the world and how we can, do, and should go about it.

This question brings to the fore two fundamental principles of knowledge: *epistemic friction* and *epistemic freedom*. Knowledge requires friction or constraints. Most importantly, knowledge is substantially constrained by its target—the world. But knowledge also requires active involvement of the theorist in both discovery and justification, i.e., the exercise of freedom. Freedom and friction, however, are not disjoint. In particular, freedom itself is an important source of constraint. Thus, the constraints of truth, justification, and substantive (informative, explanatory) theorizing are imposed on us by ourselves, through norms that, exercising epistemic freedom, we ourselves create. These normative constraints are universal, applying to knowledge qua knowledge; hence they encompass *all* fields of knowledge, from the natural and social sciences to logic, mathematics, and philosophy. It is with these two basic principles of friction and freedom that I begin the essay.

Two of the main friction requirements applicable to philosophy are (i) grounding our system of knowledge (the compilation of our most advanced theories) in the world, and (ii) substantive theorizing. By a philosophical “grounding” I mean a critical explanatory and normative theory that tackles the central epistemic question delineated above, namely, whether, to what extent, and how we humans are in principle capable of acquiring genuine knowledge of the world, and what constraints such knowledge should satisfy. Accordingly, my project is the classical project of a “foundational” theory of knowledge. Today, however, the classical foundational project is widely viewed as a dead project. This is not surprising. Traditionally, the foundational project was associated with a highly problematic—arguably self-defeating—methodology, the so-called “foundationalist” methodology, and the failure of this methodology is viewed by many as a failure of the foundational project itself. In my view, this identification is unwarranted. The foundationalist methodology is just one methodology for pursuing the foundational project; nothing prevents us from devising a new foundational methodology that avoids its pitfalls.

My first task, therefore, is to develop a methodology for a “foundation without foundationalism”¹. The foundational methodology I develop in Part I differs from its foundationalist predecessor in being *holistic*, and from other new foundational methodologies in its emphasis and scope. I call it “foundational holism”. Instead of using rigid tools and demanding a strict ordering of the grounding process, it offers flexible, holistic tools, setting no preconceived demands either on the order in which the grounding is conducted or on the resources used. And instead of having a partly coherentist orientation and being limited to empirical knowledge, it is robustly world-oriented and universal. Furthermore, its conception of grounding is not “thin” or watered-down. Indeed, in certain significant ways foundational holism is

¹ A term adapted from Shapiro 1991.

more demanding than foundationalism. While most foundationalist methodologies give logic a “free pass” (as far as grounding it in the world is concerned), foundational holism does not. As a field of knowledge, logic requires a grounding in the world as much as any other field, though it must be grounded in an appropriate facet of the world. We may say that foundational holism puts the holistic method in service of a robust, world-oriented, universal foundational project.

It is important to clarify that foundational holism is not a holistic methodology in the sense that the smallest unit of knowledge is our body of knowledge *as a whole*. Instead, it regards our system of knowledge as a *structured network* of relatively independent units, interrelated yet preserving their own identity. Accordingly, it takes into account both similarities and differences in the conditions under which the knowledge provided by different branches of knowledge is acquired and justified, and it grants epistemologists the freedom to ground different units of knowledge in ways that reflect both their similarities and their differences.

My next step (Part II) is the development of a general model of knowledge that exemplifies the principles of epistemic friction and freedom as well as those of foundational holism. Among the distinctive characteristics of this model, in addition to holism, are its dynamic structure, its broad conception of reality, and its view of intellect and sensory perception as equal players in knowledge. My starting point is Quine’s model in “Two Dogmas of Empiricism”. Quine’s model has a significant interface with both reality and mind (periphery and center); it is a rich holistic model, with an elaborate network of connections between diverse units of knowledge; and it rejects the traditional divisions of units of knowledge into those grounded in reality and those grounded *solely* in the mind. But while I adopt these elements of Quine’s model, I renounce others.

One of my criticisms of Quine’s model focuses on its overly static structure. Although Quine’s model is more flexible than most traditional models, its structure is still exceedingly static and its two poles—periphery and center—are exceedingly narrow. Specifically, its periphery is limited to observational sentences and its center to pragmatic considerations and conventional postulations. As a result, highly abstract disciplines, such as mathematics and logic, are barred from the periphery. Untangling the metaphor, logic and mathematics, in Quine’s model, are not subject to exacting standards of veridicality; as such, they are largely frictionless. Logic and mathematics can be challenged obliquely by the world, through observational statements, but this challenge is limited to those parts of their content that impact observational statements. There is no room in Quine’s model for a veridical challenge to logic and mathematics as branches of knowledge in their own right.

This limited view of abstract knowledge is closely connected to Quine’s deep, and quite radical, empiricism—another target of my criticism. Not only is Quine’s conception of center and periphery relatively narrow, but so is his conception of mind and world. Quine never asks whether objects in the world have abstract features, nor does he ask whether humans have resources for cognizing such features.

To me, it is inconceivable that humans would reach the level of knowledge they have without a significant contribution of intellect (as something distinct both from sensory perception and from pragmatic conventions), but Quine completely neglects the role of intellect in knowledge. He considers, and rightly rejects, supernatural means of discovery and justification, but human intellect is not supernatural. Human intellect cannot be identified with either telepathy or clairvoyance, nor is it related to Greek deities, or the like. In the entire Quinean corpus there is no consideration of intellect as a crucially, or even potentially, significant cognitive resource.

My own model differs from Quine's on both these counts. First, it is a dynamic model: center and periphery are job descriptions rather than fixed locations, and each discipline moves from periphery to center and vice versa according to the task at hand. In this way, observational science is subject to conceptual, linguistic, and pragmatic norms; logic and mathematics are subject to veridical norms. Second, both my center and periphery are broader than Quine's, representing a broader conception of world and mind. Objects in the world, in my model, have both abstract and concrete features, and our system of knowledge aims at knowledge of both. My paradigm of an abstract feature is a formal property, for example, a cardinality property. Most properties of objects—including physical properties of physical objects—have cardinality properties (for example, the property of being a moon of Earth has cardinality ONE). These properties, like many physical properties, are governed by laws (regularities), and our epistemic goals include knowledge of those laws.

My conception of mind is also broader and more open-ended than Quine's. Far from being purely pragmatic, it includes intellectual resources of discovery and justification, working in tandem with sensory perception. My paradigm of intellectual knowledge, too, is relatively broad. I call it "figuring out". Figuring out, in the everyday sense of the word, is something that humans engage in at any age (babies constantly figure out things), in any context, and on any level of abstraction (from figuring out why a given computer crashed to figuring out whether a given logical system is complete). It is also something that goes beyond mere sense perception, involves discovery, and is not limited to pragmatic considerations. The theoretical investigation of figuring out is a job both for psychology and for philosophy. The latter job is within the scope of *Epistemic Freedom*. Here I emphasize the centrality of figuring out (or something like it) for all knowledge, abstract as well as empirical.

One of the main instruments of epistemic friction is a standard of truth, one that measures the success of theories with respect to the *correctness* of what they teach us about the world. Such a standard is required by our model of knowledge, and in Part III I develop a theory of truth that introduces such a standard. This theory is substantivist (rather than deflationist), and its standard of truth is a correspondence standard, albeit a non-traditional correspondence standard.

Today, many philosophers are attracted to the deflationist approach to truth. But in my view truth is too substantial, too complex, and too important to be adequately

approached in a deflationist manner. Having found the reasons put forth in support of deflationism unconvincing, I ask: Is there anything inherently problematic about truth, anything that stands in the way of substantive theorizing about it? My answer to this question is positive. There is something inherently problematic about truth, or rather about the way we, philosophers, approach it. On the one hand, truth applies to an enormous array of highly diverse and often complex cognitions (thoughts, sentences, theories); on the other hand, philosophers traditionally expect the theory of truth to take the form of a single and simple definition or definition schema. But there is no reason to surmise that truth applies to all cognitions in the same way or always in a simple manner. One way to deal with this problem is to overlook the complexity of truth so as to preserve a neatly unified conception of its theory. This way leads to a frictionless theory. Another is to recognize the complexity of truth, and in particular the tension between *unity* and *disunity* (diversity) inherent in it, and to look for ways to resolve this tension.

In our time, the tension between unity and disunity has been studied primarily in connection with science (rather than with truth). The subject matter of science—nature as a whole—is a broad and diverse subject matter, and the tension between unity and disunity poses a serious problem for scientists. But this problem is not unsolvable: “every science needs for its healthy growth a creative balance between unifiers and diversifiers” (Dyson 1988: 47). This solution has nothing to do with the empirical nature of science and as such it is applicable to all fields, including philosophy. My solution to the problem of the unity and disunity of truth follows the same line of reasoning: the theory of truth is best viewed as a system of substantive principles of varying degrees of generality. Some of these principles are unifying in nature, others are attuned to the diversity of truth. All, together, offer a “fruitful balance” between unity and diversity.

To arrive at these principles I approach truth from a cognitive rather than from a linguistic perspective. Instead of starting with the totality of true sentences and asking “What is common to all these sentences?”, I start with the basic human cognitive situation and ask “Under what conditions does truth arise as a significant standard for human cognitions?” My answer is that truth arises at the juncture of three fundamental modes of human cognition which I call “immanence”, “transcendence”, and “normativity”. In the immanent mode we direct our cognitive gaze at the world (thinking in the way one thinks when one stands *within* a theory); in the transcendent mode we move beyond our immanent cognitions to a standpoint in which we hold both these cognitions and the world in view; and in the normative mode we ask certain critical questions concerning the relation between the two. I call the principle that truth arises at the juncture of these three modes of thought “the fundamental principle of truth”. (It is worthwhile to note that the appeal to transcendence does not involve commitment to a “God’s eye view”: transcendence, here, is a move from one human standpoint to another, one which is more powerful in certain respects, yet still decidedly human.)

My second core principle of truth is a non-traditional *correspondence* principle. Like the traditional correspondence principle, it sets non-trivial demands on our cognitions with respect to their relation to the world; but unlike traditional correspondence it does not determine in advance what pattern this relation can, does, and should exhibit. In thinking about correspondence I am guided by two ideas. Given the value we place on correctness, we need a robust correspondence standard. But as a standard for creatures in a complicated cognitive situation, fraught with difficulties and requiring considerable ingenuity and creativity, it must be flexible and adjust to the ways we tackle diverse cognitive challenges. I call this type of correspondence “manifold correspondence”.

To provide an example of a non-traditional pattern falling under “manifold correspondence” I turn to mathematics. Here, too, my approach is a bit uncommon. Instead of starting with existing mathematical theories and asking what facets of the world they correspond to, I start with the world and ask: is there any aspect of the world whose knowledge requires something like a mathematical theory? And instead of using current language, which historically developed in a partly haphazard way and for multiple purposes, as a strict yardstick, I ask: how, given our cognitive resources and limitations, would we best employ the language we have to develop theories of this aspect? My answers are, first, that there *is* a particular facet of the world whose knowledge requires a mathematical discipline, namely, its *formal* facet; and second, that the linguistic expressions we use to describe this formal facet correspond to it in a “composite”, indirect way. First-order arithmetic, on this account, is a theory of finite cardinality properties rather than of numerical individuals in the world. Its ability to reach its target in the world is due, to a large degree, to the ingenious, yet disciplined, exercise of epistemic freedom by laypeople and mathematicians.²

Some of the thorniest challenges facing any foundational methodology, model of knowledge, and theory of truth are posed by logic. Logic seems to resist many of the requirements imposed on all other disciplines, and this tempts us to grant it special allowances. Logic is often released from the friction requirements of substantiveness and grounding-in-reality, from the need to face the world in the periphery, from a robust standard of truth, and so on. Instead, logic is frequently regarded as purely conceptual, true by convention, trivial, obvious, a mere game, a matter of choice of language, a matter of convenience, and so on. This is not surprising, given its fundamentality, generality, abstractness, and formality. But exempting logic from central epistemic requirements is also problematic. After all, due to its crucial—theoretical as well as practical—role in our system of knowledge, it is especially important to make sure that logic does not malfunction, does not introduce error into, and potentially compromise, the integrity of our entire system. This means that logic—as much as, and even more than, any other discipline—requires a systematic

² The route to full-fledged mathematics involves expanded principles of a similar kind.

veridical foundation. In Part IV, I delineate an outline of a theoretical foundation for logic. This foundation conforms to the universal principles of epistemic friction, explains why, and in what sense, logic has a foothold in the periphery, and subjects logic to a robust standard of truth. In a nutshell: a claim of logical consequence is true if it corresponds to (an appropriate instance of) an appropriate law governing the world; false otherwise.

Among the advantages of this account is its ability to offer a substantive explanation of the characteristic traits of logic: its considerable generality, topic neutrality, basicness (fundamentality), especially strong modal force, normative power, etc. The account explains these traits in terms of a highly informative notion: *invariance*. Roughly, these traits are due to the strong degree of invariance of those features of the world that logic is grounded in (which, in turn, are responsible for the special strength of those laws that true statements of logical consequence correspond to).

Another advantage of the account is its ability to unify logic and mathematics without reducing one to the other. This it does by providing a *common grounding* for both while distinguishing their *job descriptions*. Both logic and mathematics are grounded in the *formal facet of the world*, but while mathematics *studies* the formal, logic *develops a method of inference* based on it. The formal itself is explained in holistic, semi-Aristotelian terms. This account shares the methodological advantages of traditional logicism while avoiding one of its major (though largely overlooked) methodological shortcomings. Whereas logicism enables us to solve two difficult problems—the foundational problem of logic and the foundational problem of mathematics—by reducing them to one problem—the foundational problem of logic—it leaves logic itself without a grounding. In contrast, our account offers a joint grounding for logic and mathematics without ignoring their difference.

But the most important advantage of our account is its ability to ground, and provide a substantive explanation of, the *veridicality* of logic: explain why and how logic works in the world, under what conditions its claims are true or false, what distinguishes correct and incorrect logical theories, and so on. In this way it brings logic in line with all other disciplines, while recognizing—and explaining—its unique characteristics.

The essay ends with a forward look at *Epistemic Freedom*. Epistemic freedom is inherently complicated: without it, we are unable to achieve, or even pursue, our epistemic goals; but in exercising it we introduce risk and uncertainty into the pursuit of these goals. The present volume emphasizes the complementarity of friction and freedom, construes freedom as a central element of the basic human epistemic situation, and integrates it into its model of knowledge, its account of intellect, its theory of truth, and its logical foundation. Still, much work remains to be done. We need to study the scope and limits of epistemic freedom, investigate its tension and cooperation with epistemic friction, learn how to integrate it into successful epistemic strategies, guard against its excesses, and so on. Most importantly, we need to better understand its working, its role in discovery and justification, the cognitive faculties exercising it, and its manifestations in all fields of philosophical study.

Contents

Acknowledgments

xvii

Part I. Epistemic Friction

1. Epistemic Friction and Freedom	3
1.1 Epistemic Friction and its Roots in Kant, Wittgenstein, and McDowell	4
1.2 The Present Conception of Epistemic Friction	8
1.3 Epistemic Freedom	12
1.4 Between Friction and Freedom	16
2. A Sustainable Epistemic Methodology	17
2.1 Epistemic Friction and the Illusion of Foundationalism	17
2.2 From Foundationalism to <i>Foundational Holism</i>	20

Part II. A Dynamic Model of Knowledge

3. Quine's Model of Knowledge: An Inner Tension	39
3.1 The Initial Promise of Quine's Model	39
3.2 An Inner Problem in Quine's Model	40
3.3 Objections and Responses	43
3.4 Dummett's Solution to the Inner Problem	48
3.5 A New Solution	56
4. Dynamic Model—Two Dimensions of Change	62
4.1 Contextual Dynamic	62
4.2 Temporal-Developmental Dynamic	68
5. Reality, Intellect, Realism	73
5.1 The Basic Epistemic Situation	75
5.2 Reality: Beyond Platonism and Nominalism	77
5.3 Intellect: Beyond Apriorism and Empiricism	84
5.4 Robust yet Non-Rigid Realism	91
6. Differences with Quine	100
6.1 Contrasts with Quine's Model	100
6.2 Immunity to Criticisms of Quine's Model	112

Part III. The Structure of Truth

7. A Substantivist Theory of Truth	131
7.1 A Substantivist (as Opposed to Deflationist) Methodology	131
7.2 The Unity and Disunity of Truth: Challenges and Strategies	134

8. Basic Principles of Truth	162
8.1 The Fundamental Principle of Truth	162
8.2 Ramifications for Skepticism	175
8.3 The “Manifold” Correspondence Principle	186
8.4 Application to Mathematics (A New Theory of Mathematical Truth)	192
8.5 The Logicality Principle (Tarski’s Theory of Truth in Perspective)	218
 Part IV. An Outline of a Foundation for Logic	
9. The Foundational Problem of Logic	239
9.1 The Foundational Problem of Logic as a Methodological Problem	239
9.2 Analysis of the Problem and Alleged Remedies	243
9.3 Solution: The Foundational-Holistic Methodology	250
10. An Outline of a Foundation for Logic	253
10.1 What is Logic’s Task in our System of Knowledge?	254
10.2 Is Logic Grounded in the Mind or in the World?	255
10.3 Why does Logic Require a Grounding in the World?	260
10.4 What Specific Features of the World is Logic Grounded in?—The Formality Thesis	271
10.5 From Formality to Generality, Necessity, Topic Neutrality, Strong Normativity, Quasi-Apriority, and More	288
10.6 Source of the Normativity of Logic, Tarski’s Problem, Truth and Logical Truth, and Other Issues	294
10.7 Questions and Objections: Logical Constants, Invariance, Generality, and Necessity	302
10.8 Logic and Mathematics: An Alternative to Logicism	320
10.9 On the Possibility of Error and Revision in Logic	327
10.10 The Scope of Logic	331
Conclusion: Toward Freedom	339
 <i>References</i>	 343
<i>Index</i>	359

Acknowledgments

In writing this essay I benefited from discussions with many people. I was fortunate to visit the Hebrew University of Jerusalem and other Israeli universities a number of times, and my philosophical exchanges with colleagues, friends, and students there were extremely fruitful and enjoyable. My colleagues and students at the University of California, San Diego (UCSD) contributed to a friendly and intellectually stimulating environment. In particular, I would like to thank my students Matt Evpak, Tanya Hall, Tomoya Sato, and Sebastian Speitel, and my former student and current colleague, Cory Wright. I received numerous comments on my work and had multiple stimulating conversations with colleagues in conferences, colloquia, over email, and more. John MacFarlane, Pen Maddy, Oron Shagrir, Stewart Shapiro, and Mark Steiner, among others, offered comments that led me to add, remove, or change several parts of the work. Charles Parsons' mentorship had a long-lasting impact on my work. Philip Kitcher was very supportive of my work during my early years at UCSD, and I shared with him many of my ideas on Quine. My main philosophical interlocutor in the last decade has been Carl Posy. His stimulating ideas, criticisms, comments, and suggestions, as well as his friendship, were very important for me. Peter Momtchiloff from Oxford University Press first approached me about writing a book during a conference in London in 1998 and since then has been extremely encouraging and supportive. The anonymous readers of this essay wrote perceptive and astute reviews that helped me prepare the final version of the manuscript. Indeed, it was only after reading their reviews that I saw my way into writing a preface for the essay. Matt Evpak prepared the index and read the manuscript of the book, and I am greatly obliged to him.

My deepest gratitude, love, and appreciation go to my sons, Itai and Shlomi, to whom I dedicate this book, and to my husband, Peter. In addition to filling my life with meaning and happiness, Itai and Shlomi shared their ideas with me and commented on mine. In scores of conversations they influenced my thinking on many of the topics discussed in this essay. My husband Peter is, and has always been, my closest companion and my most important source of inspiration. Peter read every version of this work, and commented in great depth and detail on all its aspects, from content and structure to style and tone. His encouragement, sustenance, and love are invaluable to me.

I would also like to thank several institutions for fellowships that helped me work on this book: the Lady Davis Foundation and the Hebrew University of Jerusalem for a fellowship and a visiting professorship; the Center for Language, Logic, and Cognition at the Hebrew University for a visiting professorship; the University of Santiago de Compostela, Spain for a visiting professorship; the Brazilian Research

Council for a research grant; the UCSD Humanity Center for two fellowships; and the Dean of Humanities at UCSD for an Innovation Fund Award.

Finally, I am grateful to the *Journal of Philosophy* and the *Bulletin of Symbolic Logic* for permissions to include material from my publications. “Is There a Place for Philosophy in Quine’s Theory?” (*Journal of Philosophy* 96, 1999: 491–524) is reproduced in Chapters 3 and 4. Large segments of “In Search of a Substantive Theory of Truth” (*Journal of Philosophy* 101, 2004: 5–36) appear in Part III, and some segments of “The Foundational Problem of Logic” (*Bulletin of Symbolic Logic* 19, 2013: 145–98) are integrated into Part IV.

PART I

Epistemic Friction

1

Epistemic Friction and Freedom

My starting point is the observation that every rational act involves both *freedom* and *constraint*: freedom to act and set rational standards for our actions, and constraints imposed by our environment on the one hand and our standards of rationality on the other. The development of a system of knowledge—a body of disciplines and theories¹ seeking knowledge of various aspects of the world (in a broad sense of the word)—is also a rational enterprise, and as such it, too, requires both freedom and constraint. Freedom, here, is freedom to actively engage in epistemic pursuits: set up our epistemic goals, choose the subject matter of our investigations, ask questions, select and apply epistemic norms, design research programs, construct epistemic tools, do experiments, make calculations, draw conclusions, devise strategies, make practical and theoretical decisions, etc. And constraint is constraint coming from two sources, the world and the mind. The world as the object or target of our theories restricts what we can truly say about it, and the mind restricts our theories both voluntarily and involuntarily: voluntarily, through our chosen goals, standards, and decisions, and involuntarily, through our makeup and built-in limitations. We may say that neither freedom without constraint nor constraint without freedom can give rise to knowledge. Freedom alone cannot distinguish knowledge from phantasm; constraint by itself would leave us cognitively inert.

¹ Two terminological notes:

(i) *System of Knowledge*: Throughout this essay I use the notion of *system of knowledge* as a partially idealized notion indicating the collection of disciplines that constitute our integrated body of theoretical knowledge.

(ii) *Theory*: I use the notion of *theory* in a broad, everyday sense. (For example, I do not limit “theory” to “axiomatic theory”.) In addition, I sometimes use “model” (as in “model of knowledge”) as synonymous to “theory” (“theory of knowledge”), and I do not draw a sharp distinction between a theory and its parts (e.g., the present essay develops both a *single* theory of knowledge, truth, and logic, and *three* interconnected theories, each devoted to one of these topics).

1.1 Epistemic Friction and its Roots in Kant, Wittgenstein, and McDowell

To clarify the idea of epistemic friction and identify the major friction requirements, I will start with the roots of this idea in three philosophical figures: Kant, Wittgenstein, and McDowell.²

A. *Kant*. The idea of epistemic friction is naturally traced to Kant, although Kant himself never used either the term “epistemic friction” or any other term of art for it. Kant illustrates both the attraction and the futility of frictionless theorizing in his dove metaphor:³

The light dove, cleaving the air in her free flight, and feeling its **resistance**, might imagine that her flight would be still easier in **empty space** . . . (Kant 1781/7: A5/B8–9).⁴

The prototype of frictionless theorizing, according to Kant, is traditional metaphysics. Traditional metaphysics purports to provide knowledge of “things in themselves” through our conceptual faculty (“understanding”) alone. But this faculty by itself is incapable of testing our theories or providing means for correcting them. The result is “darkness and contradictions” (Kant: Aviii), dogmatism, skepticism, endless disagreements, and fruitless controversies. An early practitioner, and victim, of frictionless theorizing is Plato:

It was thus that Plato left the world of the senses, as setting too narrow limits to the understanding, and ventured out beyond it on the wings of the ideas, in the **empty space** of the pure understanding. He did not observe that with all his efforts he made no advance—meeting **no resistance** that might, as it were, serve as a support upon which he could take a stand, to which he could apply his powers, and so set his understanding in motion (Kant: A5/B9).

Kant’s diagnosis of Plato’s error, and that of traditional metaphysicians after him, is that they “resort to principles which overstep all possible empirical employment” (Kant: Aviii). The remedy to this malady, Kant says, is to anchor knowledge in

² Since two of these figures are “historical” figures, let me make a prefatory comment on the appeal to historical figures in this essay. In attributing various views to historical figures I am fully aware of the possibility of controversy. I do not purport to present the ultimately correct or the most commonly accepted interpretation in each case; I present these authors’ ideas as I, a philosophical interlocutor, understand them. The main purpose of the historical discussions is to clarify the ideas developed in the present essay, ideas whose plausibility, however, is not dependent on either the connections or the contrasts with historical figures.

³ Throughout the work, I use **boldface** within citations to indicate my own emphases.

⁴ My citations from Kant in this volume are largely based on the Kemp Smith translation, which, to my ear, captures the gist of Kant’s ideas more eloquently than other translations, though is at times less accurate. Where questions of accuracy arise, I consult the Guyer and Wood translation. In the present citation, however, I take the liberty of diverging from both translations by using the pronoun “her” in referring back to the dove. (Kemp Smith uses two genderwise incompatible pronouns—“her” and “its”; Kant himself uses “sie” and “ihr”; and Guyer and Wood use “it” twice.)

sensory experience, which, unlike pure understanding, can test our theories and guide our inquiry. Without experience human theorizing is no more than “random groping” (Kant: Bvii and xv), a futile play of imagination and concepts. What the requirement of epistemic friction amounts to for Kant is that “we must never venture with speculative reason beyond the **limits of experience**” (Kant: Bxxiv). A broader perspective, however, suggests that Kantian friction consists not of one but of two requirements: (i) human knowledge requires constraints by *experience*, and (ii) human knowledge requires constraints by the *mind*.

What these constraints amount to, however, differs from one area of knowledge to another: the natural sciences are constrained by *actual* sensible intuition (sensory perception), while mathematics is constrained by the *pure form* of sensible intuition; the natural sciences are constrained by the forms of *concepts* (“categories”), mathematics by the forms of *intuition*; natural science must validate its claims by *experiment and observation*, mathematics—by *proof and demonstration*. Furthermore, Kant’s concern about frictionless theorizing is *not equally strong* with respect to all fields of knowledge. *Logic*, in particular (that is, “pure, general, formal” logic, in Kant’s terminology), is to a certain degree exempt from the friction requirement. While logic creates friction for other sciences (all sciences must obey the laws of logic), logic itself, as a body of knowledge, is subject to few friction constraints. Logic has to satisfy the constraints of formality and generality, but, unlike the empirical sciences and mathematics, it is not subject to *major veridicality* constraints (although it is subject to relatively minor veridicality constraints, like the correct treatment of the copula). Most importantly, there is no question of error in the major logical laws. We have to guard against error in *applying* the logical laws, but (aside from a few small details) not against error in *identifying* the logical laws; i.e., we do not have to worry that our current logical theory might err in telling us what logical laws there are.

The view that a given field of knowledge is not subject to major veridicality worries, however, requires a serious justification. But Kant never offers such a justification for his attitude toward logic. One might say that logic, for him, is special: logic is analytic. But this will not do. Analytic statements, in general, *are* subject to some veridicality constraints, e.g., they have to abide by the laws of *logic*. It is only the logical laws themselves that are presumably not subject to veridicality constraints. Why Kant was not concerned with the veridicality of logic is an open question. Perhaps it was the “wisdom” of his age, perhaps his view that logic’s role was merely *negative*. Maybe he was dazzled by the “fact that since Aristotle [logic] has not required to retrace a single step” (Kant: Bviii). Perhaps he did think there was a question of veridicality in logic, but thought it concerned only our ability to see accurately what laws are built into our mind, an ability we are guaranteed to have since the mind is (so he thought) “self-illuminati[ng]” (Cavaillès 1946/60: 359), being “perfectly transparent to itself” (Longuenesse 1993/8: 75). Or maybe the possibility of error in logic did not occur to him at all. Whatever the reason, Kant did not raise

the question of veridical friction with respect to logic in any way that approaches the seriousness with which he raised this question for science and mathematics.

What about philosophy? Not traditional metaphysics which Kant charged with being frictionless, but his own brand of philosophy—transcendental philosophy. Does transcendental philosophy stand in danger of frictionless theorizing? Can it meet the friction requirement? I think Kant's answer to the first question is positive, but his answer to the second question is partly negative. Kant's lengthy attempts to justify his transcendental theses by "deductions" show that he did regard transcendental philosophy as in need of justification, and his confidence in the success of his deductions indicates that he thought such justification was possible, at least to a degree. But transcendental philosophy, according to Kant, cannot achieve the kind of friction that underwrites "real" knowledge, namely, scientific or mathematical knowledge, and its statements, therefore, cannot attain the status of *truth*. "Deductions" can prove the "right" to apply certain concepts to experience, but not the *truth* or factuality of transcendental claims. By restricting veridical friction to experience, Kant appears to deny the possibility of knowledge in the strict sense to transcendental philosophy, characterizing it as merely "*regulative*" (Kant 1781/7: A180/B222–3).⁵

B. *Wittgenstein*. The first to use the term "friction" for the idea of constraints on knowledge was, as far as I know, Wittgenstein. Wittgenstein used this term in contrasting language-game philosophy with ideal-language philosophy:

[With ideal-language philosophy, we] have got on to slippery ice where there is no **friction** and so in a certain sense the conditions are ideal, but also, just because of that, we are unable to walk. We want to walk: so we need *friction*. Back to the rough ground! (Wittgenstein 1953/58/63: §107).

In contrast:

[L]anguage-games [do not ignore] **friction** and air-resistance. The language-games are . . . set up *as objects of comparison* which are meant to throw light on the **facts** of our language (Wittgenstein: §130).

Viewed from a broader perspective, Wittgenstein, like Kant, sees epistemic friction as making different demands on different branches of knowledge. From science it demands agreement with both the sensory world and the rules of language, from logic it demands agreement only with the latter. As for philosophy, we can

⁵ But Kant is not completely clear on the epistemic status of transcendental philosophy. For example, in section VII of the Introduction to *Critique of Pure Reason* he says that a "complete system of the philosophy of pure reason . . . may be carried into execution, analytically as well as synthetically", because "what . . . constitutes [its] subject-matter is . . . the understanding . . . in respect of its *a priori* knowledge . . . [and these] *a priori* possessions of the understanding, since they have not to be sought for without, cannot remain hidden from us" (Kant A12–13/B26). (Compare with the similar thought concerning logic mentioned earlier.)

distinguish two trends in Wittgenstein's thought. To the extent that philosophy is *descriptive*, i.e., has the task of describing the rules of actual language games, it is constrained by observed linguistic facts. But to the extent that philosophy is *therapeutic* (rather than theoretical), it is exempt from the requirement of veridical friction altogether (though not from the requirement of effectively curing our philosophical angst).

C. McDowell. More recently, the term "friction" appears in McDowell's *Mind and World* (1994). Although McDowell oftentimes talks of friction as a requisite for *content*—"We aim at ensuring **friction**, which is required for **genuine content**" (McDowell 1994: 18)—his discussion extends to knowledge as well. Relating to Kant's account of "spontaneity" (faculty of conceptual synthesis or creativity) he says:

It must be possible for justifications of judgements to include pointing out at features of the world... [so] our exercises of spontaneity [do not] run without **friction** (McDowell: 39).

And like Kant, his *explicit* concern with friction is with external constraints by the world:

We need to conceive... spontaneity as subject to control from **outside** our thinking, on pain of representing the operations of spontaneity as a **frictionless** spinning in a void... [We] crave... for **external friction** in our picture of spontaneity (McDowell: 11).

As for the nature of constraints by the world, McDowell, like Kant and Wittgenstein, seems to view such constraints as *experiential* in nature. For example, in explaining the Kantian claim that theories without content are empty—a claim he (rightly) associates with Kant's friction requirement, and one he endorses—McDowell says:

Thoughts without content... would be a play of concepts without any connection with intuitions, that is, bits of **experiential** intake. It is their connection with **experiential** intake that supplies the content, the substance, that thoughts would otherwise lack (McDowell: 4).

McDowell's main concern, however, is not with friction itself, but with a trap we are in danger of falling into when we try to *establish* friction, namely, the "myth of the given" (Sellars 1956). In seeking to ground our theories in reality, McDowell says, we are tempted to postulate the existence of a brute "given":

[We] appeal to the Given, in the sense of bare presences that are supposed to constitute the ultimate grounds of empirical judgements (McDowell 1994: 24).

But appealing to the given is of no use, since being under the control of an external force does not amount to justification; at best it amounts to exculpation:

The trouble about the Myth of the Given is that it offers us at best exculpations where we wanted justifications... If [the] impingements [of the given on our senses] are conceived as outside the scope of spontaneity [i.e., concepts],... then the best they can yield is that we

cannot be blamed for believing whatever they lead us to believe, not that we are justified in believing it (McDowell: 13).

McDowell's solution to the problem of the given is largely Sellarsian: receptivity does not make an independent contribution to the cooperation between intuitions and concepts; rather, intuitions already contain a conceptual element. We should, therefore, understand Kant as saying that human cognition is constrained by our conceptual apparatus at every level, including the level of elementary sense perception:

We should understand what Kant calls "intuition"—experiential intake—not as a bare getting of an extra-conceptual Given, but as a kind of occurrence or state that already has conceptual content (McDowell: 9).

Although McDowell himself does not put it this way, we can view him as following in Kant's footsteps by indirectly advocating a second friction requirement, namely: scientific knowledge must be grounded not just in sensory perception but also in concepts, not just in experience but also in the *mind*. McDowell's idea of groundedness in the mind, however, is different from Kant's. Groundedness in the mind is groundedness in an *inculcated* system of concepts and rules of rationality rather than in a *built-in* system of transcendental principles. Another difference between McDowell and Kant has to do with the scope of conceptual friction. Unlike Kant, McDowell insists that empirical knowledge is grounded in concepts all the way down to the most elementary experiential intake.

As for logic, mathematics, and philosophy, McDowell says nothing about friction in any of these disciplines. But while it is not clear what his view of veridical friction in logic or mathematics is, his advocacy of a quietist approach to philosophy suggests that he exempts philosophy from the friction requirement, or at least its veridicality component.⁶

1.2 The Present Conception of Epistemic Friction

My own conception of epistemic friction follows Kant's, Wittgenstein's, and McDowell's on some points, and diverges on others. From Kant I take the basic idea of epistemic friction, the need for a dual grounding of knowledge in the world and the mind, the importance of providing a foundation for knowledge, the interest in unity and systematicity, and the view (reflected in his practice) of philosophy as a substantive discipline. From Wittgenstein I take the realization that we have to carry Kant's concerns beyond the areas to which he himself directed them (for example, to logic). And from McDowell I draw lessons about how to deal with the problem of epistemic friction (along with other classical philosophical problems) in a contemporary

⁶ MacFarlane (2004) compellingly criticizes McDowell for not applying his friction requirement to mathematics.

setting. But my approach to these authors is also critical. Specifically, I believe their construal of the problem of epistemic friction is exceedingly narrow, and as a result, their solutions are too narrow as well.

To present my own conception, I will formulate seven principles of epistemic friction accompanied by brief explanations and motivating remarks. The first two express what, I think, is right in the earlier positions; the last five attempt to transcend their limitations:⁷

FIRST PRINCIPLE

Following Kant, Wittgenstein, and McDowell, I formulate what can be called “the general problem of friction” as follows:

- (1) *The general problem of epistemic friction is the problem of setting adequate constraints on our system of knowledge so as to avoid empty theories and maximize genuine knowledge.*

The problem is threefold—a problem of design, understanding/explanation, and justification. The problem of design is the problem of figuring out how to design our system of knowledge so it satisfies appropriate constraints. The problem of understanding/explanation is the problem of identifying the main mechanisms of epistemic friction and showing how they create the appropriate constraints. And the problem of justification is the problem of justifying our design in light of these constraints. These problems are interrelated: our design depends on our understanding and requires justification.⁸

SECOND PRINCIPLE

A central aspect of epistemic friction, emphasized (directly or indirectly) by Kant, Wittgenstein, and McDowell, is the importance of groundedness and its two facets: groundedness in reality and groundedness in the mind. Accordingly:

- (2) *A central constituent of epistemic friction is the requirement that our system of knowledge be well grounded. This requirement has two parts: (a) our system of knowledge must be grounded in reality or the world, and (b) our system of knowledge must be grounded in the mind. Groundedness in the world is veridicality, i.e., compliance with substantial standards of truth, evidence, and justification. Grounding in the mind might include conformity with transcendental principles (Kant), compliance with inculcated principles of rationality (McDowell), and agreement with rules of language (Wittgenstein).*

⁷ Inevitably, these brief remarks will not do justice to the richness and subtlety of Kant’s, Wittgenstein’s, and McDowell’s views. But to enter into a lengthy discussion of their views would divert us from our main task, which is constructive rather than interpretive.

⁸ For a partially similar, twofold conception of “the problem of epistemology” see Sosa (1983).

But it also importantly includes other things as noted in Principles (5) and (6) below.

Clarifications: I use “world” and “reality” synonymously here (as well as throughout the volume). Also, at this initial stage I prefer to leave the boundaries of world/reality an open question. But the notions of *reality* and *system of knowledge* are coordinated in the sense that the main targets (direct or indirect) of all branches of knowledge fall within the boundaries of reality. As a result, extending the scope of our conception of *knowledge* involves extending the scope of our conception of *reality* and limiting the scope of our conception of *reality* means limiting the scope of our *system of knowledge*.

THIRD PRINCIPLE

Kant, the later Wittgenstein, and McDowell identify the grounding of knowledge in the world with its grounding in *experience*. This identification, however, is too narrow. The claim that knowledge requires friction holds for any conception of knowledge, be it empiricist or non-empiricist; therefore, the grounding-in-reality requirement should be formulated, at least at the outset, more broadly, as applying to any type of knowledge. This concurs with our decision to leave the scope of reality open at this stage. Accordingly:

- (3) *The grounding of knowledge in the world is not necessarily restricted to experience. Knowledge has to be grounded in reality, and while experience undoubtedly plays an important role in its grounding, the nature of the grounding reality and the grounding mechanisms is an open question. In particular, the possibility that knowledge is partly grounded in non-experiential facets of reality, or that the grounding mechanisms are partly non-experiential, is not ruled out.*

FOURTH PRINCIPLE

Kant, Wittgenstein, and to a limited degree McDowell think of the grounding of knowledge in the mind as primarily a grounding in *non-voluntary* principles of cognition, language, and rationality. Kant, for example, thinks of human knowledge as grounded in certain pure forms of “intuition” and “understanding”, where these forms are built into human cognition and can in no way be affected by any choice or decision humans make. (Although he talks about “spontaneity”, spontaneity, as he thinks of it, is not a voluntary activity.) Thus, if causality, say, is built into our understanding as a basic concept or form of synthesis, then nothing we discover or realize can change the requirement that our scientific theories couch their findings in causal terms. But this conception is too narrow: although knowledge is partly grounded in involuntary cognitive principles of various kinds, it is just as importantly grounded in *deliberate* cognitive principles, i.e., principles *freely, voluntarily, and actively* chosen by us. Indeed, many cases of our own involuntary conditioning are

open to critical evaluation by us, which is voluntary in nature, and to consequent refusal to abide by them, which is also largely voluntary. (Thus, we may decide that, given the development of modern science, its grounding in reality need not be causal, that some linguistic conventions should not be seriously entertained because they hamper our search for truth, that some principles of rationality inculcated in us by our teachers or parents are wrong headed, etc.)⁹ Accordingly:

- (4) *The grounding of knowledge in the mind has two aspects, passive and active. Knowledge is grounded in built-in or naturally developed principles of cognition; but it is also grounded in principles freely, intentionally, and critically developed, selected, and/or decided upon by us.*

FIFTH PRINCIPLE

Kant, Wittgenstein, and McDowell each emphasize one aspect of the mind—transcendental, linguistic-conventional, and *Bildung* (inculcation), respectively. But in fact knowledge is grounded in the mind in a far broader sense, both on the voluntary and on the involuntary level:

- (5) *Knowledge is grounded in the mind in a broad sense, both on the voluntary and on the involuntary level. On the voluntary level it is grounded in methodological, theoretical, and pragmatic-practical principles, from unity, systematicity, and veridicality (thought of as a norm created by the mind) to simplicity, economy, and practical applicability. On the involuntary level it is grounded in physical, biological, psychological, and possibly transcendental principles. And on a mixed, partly involuntary, partly voluntary level, it is grounded in social, political, conceptual, linguistic, and other principles.*

SIXTH PRINCIPLE

An important friction requirement, overlooked (as a friction requirement) by my predecessors, is *substantiveness*. It is true that one of the most important reasons a theory might hover idly in empty space is its refusal to be constrained by reality (or its authors' refusal to subject it to constraints by reality), but a no less important reason is *emptiness* or *triviality*. An empty or trivial theory cannot conflict with anything; as such it is not subject to significant constraints. To avoid lack of friction due to triviality, it is required that all branches of knowledge be substantive. Accordingly, an important constraint by the mind is *substantiveness*. Today, the substantiveness requirement is especially important in philosophy, since the popular trends of

⁹ McDowell seems to hold an intermediate position. On the one hand he regards our standards of rationality as grounded primarily in our *Bildung* (over which we have no active control), on the other hand he does allow an active critical attitude toward these standards.

deflationism, minimalism, and quietism threaten to promote trivial theories.¹⁰ We thus add:

- (6) *A central friction requirement is substantiveness: the setting of high standards of discovery, explanation, justification, informativeness, depth, theoretical significance, rigor, systematicity, intellectual interest, and the like for all theories in our system of knowledge.*

SEVENTH PRINCIPLE

While Kant, Wittgenstein, and McDowell restrict the scope of the friction requirement (either explicitly or implicitly) to certain branches of knowledge, I believe this requirement applies to knowledge *qua* knowledge, i.e., to *all* branches of knowledge. These include logic, philosophy, and mathematics along with theoretical and experimental science. Take *logic*, for example; since (contrary to prevalent opinion) logic, as a discipline, has a particular subject matter, namely *logical validity*, *logical consistency*, etc. (see Chapter 10, Section 1), and since its job is to develop a theory or theories about this subject matter, we must subject it to veridicality constraints. Thus, when a logical theory says that an inference, \mathfrak{F} , is logically valid, this claim has to satisfy several *factual* requirements: (i) that the truth of \mathfrak{F} 's premises in fact *guarantees* the truth of its conclusion, (ii) that this guarantee is in fact *modally robust*, (iii) that this guarantee is in fact *due to the logical structure* or logical content of \mathfrak{F} (rather than to its non-logical structure or content), etc. That is to say:

- (7) *The central friction requirements—groundedness in reality (veridicality), groundedness in the mind, substantiveness, etc.—are universal. They apply to, and are in principle satisfied by, our system of knowledge as a whole and each of its branches individually. This includes logic, philosophy, and mathematics, along with the social and natural sciences.*

These seven principles of epistemic friction are not intended as final. They likely require some fine-tuning and perhaps significant revision. But their fine-tuning and revision would have to satisfy principles of the same kind. In this essay I will focus on three universal constraints highlighted by the current principles: *grounding-in-reality*, *unity*, and *substantiveness*.

1.3 Epistemic Freedom

Epistemic friction is complemented by another overarching principle, *epistemic freedom*. Our notion of epistemic freedom includes both elements that one would naturally place under this notion and elements that are not naturally placed under it.

¹⁰ While the claim that philosophical theories (or some philosophical theories) should be trivial is itself not trivial, it promotes trivial theorizing in philosophy.

The underlying idea is that many factors play a significant role in the construction and success of a system of knowledge. These include both the world as target of our theories and other factors that are independent of the world. The latter include reflection, choice, decision, and deliberate action, which we normally regard as falling under “freedom”. But they also include things which are normally not subsumed under this notion, for example, our cognitive capacities and limitations. To include all these factors under the same category we *extend* our notion of *epistemic freedom* to all factors that are metaphorically *free from the world (as the target of a given theory)* in the sense of being independent of, or distinguished from it. If we regard the *realm of mind* as encompassing all epistemic factors that do not belong to the world (as target), we can say that the *realm of epistemic freedom* is the realm of mind.

Our notion of *epistemic freedom* can, thus, be viewed as a composite of two notions:

(a) *Epistemic freedom in the “natural” sense*: The construction of a system of knowledge is to a large degree an intentional project, involving a wide array of deliberate, voluntary choices and actions, including asking questions, setting goals, selecting and imposing norms, choosing and using methods and methodologies, designing, conjecturing, deriving, calculating, inventing, reasoning, solving problems, abstracting, generalizing, observing, conducting experiments, making revisions, and so on.

(b) *Epistemic freedom in an “extended” sense*: While knowledge is affected by features of the world as its subject matter, it is also affected by other factors: the structure of our mind, our cognitive capacities, social and environmental factors, human norms and methodologies, human decisions and actions, and so on. All these factors belong in the scope of our extended notion of epistemic freedom.

The two notions are not disjoint: (a) is included in (b), (b) extends (a). Nor is epistemic freedom as a whole disjoint from epistemic friction. Many products of epistemic freedom are constraining in nature, and many types of epistemic friction require freedom. Norms, for example, stand in the intersection of freedom and friction. Norms are *products* of freedom but *instruments* of constraint: they emanate from us rather than from the world, but they constrain us in developing theories of the world. In a sense, it is because of our epistemic freedom that we need epistemic constraint: as agents free to create ideas we need constraints on our creative activities in order to produce *knowledge*.

Freedom, however, goes beyond *constraint*. An important part of freedom is *enablement*. Freedom as constraint is *negative*, freedom as enablement—*positive*. Both can be either *passive* or *active*. Imposing normative standards on our practices is active, being constrained by our biological makeup is passive. Being enabled by natural and inculcated cognitive capacities is passive; deliberately creating tools, choosing norms, deciding on courses of action, etc. is active. Constraint, too, may be indirectly enabling. Constraints limit what we can (or are licensed to) do, yet by setting these limits they map out a space for what we are able (allowed) to do.

Many philosophers have written about epistemic freedom, though not under this caption, and different philosophers have focused on different aspects of such freedom. Consider, for example, Kant, Quine, and Popper. Kant concentrates on the transcendental aspects of epistemic freedom. The cognizer is a being whose two enabling faculties are intuition and understanding. Intuition contributes to knowledge through perception (both sensory and apriori perception); understanding contributes to knowledge through conceptual synthesis of representations (both perceptual and conceptual representations). Kant characterizes the cognizer's participation in the process of knowledge through the understanding as "spontaneous", and spontaneity he characterizes as "active". But while Kantian spontaneity is a form of freedom in our sense, it is a form of *passive* freedom. As we have noted before, neither the pure forms of intuition nor the pure forms of the understanding are under our active control.¹¹

Quine concentrates on different aspects of epistemic freedom: its naturalistic and pragmatic aspects. Following Carnap, Quine attributes great importance to pragmatic considerations in shaping our science (see Chapter 3 and Chapter 6, Section 1). And as a naturalist, Quine characterizes human knowledge as the outcome of natural processes occurring in a physical object:

[A] natural phenomenon, viz., a physical human subject... is accorded a certain... input—certain patterns of irradiation in assorted frequencies, for instance—and in the fullness of time the subject delivers as output a description of the three-dimensional external world and its history (Quine 1969a: 82–3).

The ensuing corpus of knowledge has two sources: an external source, providing sensory input, and an internal source, activating cognitive functions over this input. As a naturalist, Quine characterizes epistemic freedom in behavioral, physiological, and neurological terms, but other naturalists describe it other terms as well, e.g., evolutionary, social, and (non-behavioral) psychological terms. Whole disciplines—cognitive science, cognitive psychology, sociology of science, branches of linguistics, and branches of neuroscience—are devoted to a naturalist study of epistemic freedom.

Popper, in turn, emphasizes the importance of *active* freedom. Speaking of empirical science, he says:

[E]xperiment is planned action... We do not stumble upon our experiences, nor do we let them flow over us like a stream. Rather, we have to be **active**: we have to '*make*' our experiences. It is we who always **formulate** the questions to be put to nature; it is we who **try** again and again to **put** these questions so as to **elicit** a clear-cut 'yes' or 'no' (for nature does not give an answer unless **pressed** for it). And in the end, it is again we who **give** the answer; it is we ourselves who, after severe scrutiny, **decide** upon the answer to the question which we **put** to nature (Popper 1959: 280).

¹¹ For recent discussions of Kant's notion of spontaneity see e.g., Hanna (2004/13) and Tolley (2006).

Active acts of freedom are emphasized by working scientists as well. Take, for example, Watson's memoir, *The Double Helix* (1968). Watson describes the discovery of the DNA structure as achieved neither by following a rigid algorithm nor by experiencing a sudden flash of inspiration. Rather, the discovery process unfolds in a series of *decisions*, *knowingly* made at each step of the way; for example:

[A] **decision would have to be made** whether the chains would be held together by hydrogen bonds or by salt linkages involving the negatively charged phosphate groups (Watson 1968: 39).

Decisions had to be made about the number of polynucleotide chains within the DNA molecule (Watson: 55).

[W]e had **decided** upon models in which the sugar-phosphate backbone was in the center of the molecule (Watson: 56).

By now I had **decided** to mark time by working on tobacco mosaic virus (TMV) (Watson: 75). I **decided** to risk a full explosion (Watson: 106).

I tried to **decide** between the two- and three-chain models. . . . I had **decided to build** two-chain models (Watson: 108).

I **decided** that no harm could come from spending a few days building backbone-out models (Watson: 112).

Logic, too, supplies striking examples of active freedom. An extraordinary example of an act of freedom in logic is Frege's creation of a new symbolic language for the expression of logical inference, *Begriffsschrift*. The creation of this language is considered by many the single most important event in the history of logic since Aristotle.

A prime example of active freedom in philosophy is Kant's "Copernican" revolution. Having become disillusioned with the traditional approach to justifying our claims to knowledge, Kant deliberately set out to change it by changing the basic "gestalt" of our cognitive situation. Instead of placing the world "in itself" in the center and us, prospective knowers, as circling around it and trying to fathom it from a standpoint external to it, he decided to turn the traditional picture around and view ourselves as standing in the center while the world is circling around us. This turnabout, Kant argues, enables us to establish the validity of science in a way that defies Humean skepticism. Leaving aside the question of how successful Kant's revolution in fact was, there is no doubt that it changed the face of philosophy.

Active freedom in knowledge, however, is not limited to groundbreaking acts like those of Kant, Frege, or Crick and Watson. Active freedom is required in every act of rationality, groundbreaking or quotidian. The mutual interdependence of active freedom and rationality is explained by Sen as follows:

Even though the idea of freedom is sometimes formulated independently of . . . reasons, freedom cannot be fully appraised without some idea of what a person . . . has a reason to prefer. Thus, there is a basic use of rational assessment in appraising freedom, and in this sense, freedom must depend on reasoned assessment of having different options. . . . Rationality as the use of reasoned scrutiny cannot but be central to the idea and assessment of freedom.

[R]ationality, in its turn, depends on freedom. This is not merely because without some freedom of choice, the idea of rational choice would be quite vacuous, but also because the concept of rationality must accommodate the diversity of reasons that may sensibly motivate choice. (Sen 2002: 5).

Epistemic freedom and friction are fundamental, complementary requisites of knowledge as a rational enterprise. The “problem” of epistemic freedom is the problem of how to optimize our exercise of freedom in developing a system of knowledge.

1.4 Between Friction and Freedom

In pursuing a unified account of epistemic friction and freedom, we have to decide what to do first. This is largely a practical question. Two useful guidelines are: (a) start with problems you can handle best at the moment, given your present resources; (b) start with questions for which fewer answers are available, so you don’t tie your hands with choices you do not have to make. Based on these guidelines, I start with epistemic friction and leave the more detailed study of epistemic freedom for another essay. In studying epistemic friction I focus on six challenges: (1) Find, or develop, an epistemic methodology for grounding all branches of knowledge (including abstract branches such as logic) in reality, given the known difficulties. (2) Using this methodology, construct, or develop, a general model of knowledge that incorporates the general principles of epistemic friction and freedom, including substantiveness. (3) Show how philosophical theories within the model (such as the theory of truth) can satisfy the substantiveness requirement. (4) Construct a theory of truth for the model. (5) Delineate a foundation for logic in accordance with, and using the resources provided by, (1) through (4). (6) Aim at maximal unity. Let us turn to the first methodological task.

2

A Sustainable Epistemic Methodology

2.1 Epistemic Friction and the Illusion of Foundationalism

While the term “epistemic friction” is new and its roots relatively recent, the general problem of epistemic friction has been with us for a long time, and one of the philosophical projects associated with it is the *foundational project*, the project of a theoretical foundation for knowledge. By a philosophical foundation or *grounding* for knowledge I mean, in this essay, a substantive, largely normative, philosophical theory that investigates critically the possibility of human knowledge and studies the basic features of our system of knowledge. A foundational theory of this kind seeks to understand the overall structure of our system of knowledge, its tasks and goals, the relation between its different branches, and other relevant issues. These include (i) the ways the system as a whole and its branches individually are constrained/warranted by the mind on the one hand and the world on the other, (ii) the veridicality—truth, evidence, and justification (with respect to correctness)—of its theories, (iii) its strategies for reaching different facets of reality, (iv) its standards of accepting (rejecting, revising) different theories, and so on. The list is in principle open-ended since new interests and concerns may be raised at any time. In addition, the investigation itself is likely to raise new questions which require new foundational answers.

Today, however, the foundational project in epistemology is held in disrepute. The failure (or perceived failure) of the great foundational systems of Descartes, Kant, Carnap, and others, has left philosophers disillusioned, and the project itself is held responsible. But this attitude is based on a misunderstanding. It is not the foundational project itself which is misguided; it is our own identification of this project with a flawed methodology, *foundationalism*, which is ill-conceived. It is true that, for a long time, the foundationalist strategy has been the only available strategy for pursuing the foundational project, and this has understandably led philosophers to identify the strategy with the project. But the foundationalist strategy is but one strategy, and today we have tools for developing alternative foundational strategies or methodologies. Our first task, therefore, is to revisit the traditional foundational methodology and identify its problems in a way that is conducive to a fruitful, constructive solution.

THE FOUNDATIONALIST METHODOLOGY

This methodology seeks to establish all human knowledge on a dual foundation of (i) *basic knowledge*, and (ii) *knowledge-extending procedures*. Foundationalists are divided into several subgroups—e.g., rationalists and empiricists, externalists and internalists—but they are all united by their fundamental principles:

- (a) Our system of knowledge is divided into items of two types, *basic* and *derivable*.
- (b) *Basic items of knowledge* are grounded in reality (or whatever else they might be grounded in) directly, through direct experience, rational intuition, convention, etc.
- (c) *Derivative items of knowledge* are grounded indirectly, through reliable knowledge-extending procedures: deductive, inductive, and possibly others.

A salient foundationalist requirement, directly associated with these principles, is a *strict ordering requirement* imposed (either implicitly or explicitly) on the grounding relation and, consequently, on our system of knowledge as a whole.¹ This requirement, reflected in such familiar metaphors as the *tree* and the *pyramid*, sets formal conditions on the grounding relation. Essentially, this relation is irreflexive, antisymmetric and transitive, has a relatively small, absolute base consisting of “minimal” (“initial”, “atomic”) elements, and connects each non-minimal element to its atomic anchors in a finite number of steps.² This feature of foundationalist epistemology is both a source of its strength and a source of its weakness, both a cause of its considerable attraction and a cause of its failure. On the one hand, foundationalism reduces the unmanageable task of providing a grounding for our entire system of knowledge to the (seemingly) manageable task of grounding only its basic constituents; on the other hand, foundationalism has no resources for grounding the basic constituents of knowledge. Instead of having to establish the credentials of each and every unit of knowledge from scratch, foundationalism lets us establish only the credentials of a small number of units and rules from scratch (or at least a small number of kinds of units/rules). But ironically, the same strong ordering that makes foundationalism so promising creates formidable obstacles to the realization of its promise. Since the basic units of knowledge lie at the bottom of the foundational hierarchy, no unit (or combination of units) is suitably situated to produce the requisite resources for grounding the basic units.

¹ The observation that foundationalism is importantly characterized by “formal” requirements was made earlier by Sosa (1980a, 1980b). Sosa distinguishes between “formalist” and “substantive” foundationalism. The foundationalism I relate to here is both formalist and substantive in Sosa’s terminology.

² Clarifications: (i) The smallness of the base relates either to the number of items or to the number of types of items in the base. (ii) The grounding relation may be extended by adding definitions and equivalences which are symmetric; but in its core it is (or is reducible to) a strong and quite rigid partial-ordering relation.

This is the “*basic-knowledge predicament*” of foundationalist epistemology: Foundationalism’s success in grounding our system of knowledge depends on its success in grounding the basic units, but due to the strict ordering it imposes on the grounding relation, it has no resources for grounding those units. This predicament can be described as a breakdown of friction. Aiming at perfect friction, foundationalism ends up losing friction altogether. The problem affects all aspects of the grounding project: prevention and correction of error, account and explanation of knowledge, justification of claims for knowledge, and preparation of conceptual tools for these tasks. We may present it in argument form as follows:

Three interrelated principles of foundationalist epistemology are

1. The grounding of our system of knowledge is reduced to the grounding of the basic units.
2. To ground X we can only use resources more basic than those generated by X.
3. No unit of knowledge (or combination of such units) can generate more basic resources than those generated by the basic units.

It follows from these principles that:

4. No unit of knowledge, or combination of such units, can produce resources for grounding the basic units.

Foundationalism’s problems are sometimes blamed on its “absolute certainty” requirement. But while this requirement is indeed highly problematic, relaxing it will not, by itself, save it from the basic-knowledge predicament. To provide a justification (grounding) for the basic units, fallible or infallible, we need to prepare some cognitive resources, but the strict ordering requirement says that such resources cannot be prepared. This predicament leaves the foundationalist with two natural options:

- (a) Show that there is no need to ground the basic units.
- (b) Show that the basic units can be grounded without using any cognitive resources produced by, or with help from, our system of knowledge.

Neither option, however, succeeds in recovering the lost friction.

(a) *Foundation without a grounding of basic units.* One way the foundationalist might try to justify her approach is by arguing that it is impossible to provide a foundation for the basic units regardless of how you go about it, and you cannot require a methodology to accomplish the impossible. But no matter how this argument is worked out and whether its impossibility claim is correct, it cannot save foundationalism from the basic-knowledge predicament. On the contrary, if sound, this argument is a *reductio ad absurdum* of the foundationalist doctrine. If, to abide by its own rules, a doctrine has to accomplish the impossible, then it is an unviable doctrine that must be renounced or at least radically changed.

Another argument the foundationalist might attempt is: no matter what methodology one chooses, grounding (explaining, justifying) must stop at some point; why not at the basic units? This argument, too, cannot vindicate foundationalism. Granting the practical necessity of stopping the grounding process at some point, not all points are equal. Leaving a remote area high up in the foundationalist hierarchy ungrounded will have few ramifications for the overall grounding of our system of knowledge, but leaving the base ungrounded will subvert the entire grounding project. It is a *structural* predicament of foundationalism that leaving the higher units of knowledge ungrounded would make it redundant, while leaving the basic units ungrounded would undermine its integrity. The integrity of our system of knowledge, in foundationalist systems, comes from below, from theories in the lowest echelon of the foundational ordering; hence, by foundationalism's own lights, failure to secure the integrity of these theories undermines the integrity of the entire system.

(b) *Foundation without resources produced by our system of knowledge.* This appears to be the foundationalists' solution of choice to the basic-knowledge predicament. It is inherent in the foundationalist method, its adherents say, that the foundation of the basic units differs in kind from that of the other units: to provide a foundation for the non-basic units we require knowledge-based resources, but to provide a foundation for the basic units we do not. Their foundation is "for free", so to speak. Four contenders for a free-standing foundation of basic knowledge are: pure sensory perception, pure abstract intuition, common-sense obviousness, and conventionality. The foundational credentials of all four, however, have been widely contested. We are all familiar with Sellars's (1956) criticism of pure sensory perception as an independent basis for knowledge (under the rubric of "myth of the given") and with Hume's skeptical argument, which in one of its forms says that sensory perception, being highly particular, cannot serve as a basis for theoretical knowledge, whose laws are highly general. We are also familiar with Benacerraf's (1973) and others' criticisms of the epistemic credentials of pure abstract intuition. And, of course, Quine's (1935, 1954a) criticisms of conventions as a basis for knowledge are well known. Finally, I myself have argued against common-sense obviousness as a basis for knowledge (Sher 1999b), and I will do so again in Chapter 9, Section 2, where I will also argue against pure conventions and Platonic intuition as a basis for logic.

Of course, one cannot rule out altogether the possibility that foundationalism will overcome its impediments by internal revision. But I think that the severity of its problems on the one hand and our ability to identify a tangible source of these problems on the other suggest that a search for a new foundational methodology (one that does not involve the rigid ordering requirement) is more promising.

2.2 From Foundationalism to *Foundational Holism*

The failure of traditional foundationalism (along with related problems in traditional philosophy) has led many twentieth-century philosophers to draw quite

radical conclusions. Some, identifying philosophy with epistemology and epistemology with foundationalism, have gone so far as to reject the philosophical enterprise altogether (Rorty); others have rejected the enterprise of *theoretical* philosophy (the later Wittgenstein); still others have advocated a purely descriptive, rather than normative, epistemology (Quine in some of his modes); quite a few have relinquished, or have severely curtailed, the objective of a substantive, explanatory philosophy (all quietists and some deflationists); several have given up the goal of grounding knowledge in *reality* (traditional coherentists); and most have identified the renouncement of foundationalism with the renouncement of the foundational project itself.

These views, however, are too extreme. Epistemology is just one branch of philosophy, and the *foundationalist* methodology is just one epistemic methodology—indeed, just one *foundational* methodology. It is true that for a long time this methodology was the only existent foundational methodology, and this may explain why so many of its features have become entangled in our conception of a foundation. But this entanglement can and ought to be unraveled. To some extent it already has been; yet its unraveling has often been accompanied by giving up the foundational project itself. Foundationalism's failure, however, opens the door to a different foundational methodology, a methodology that, while free from foundationalism's encumbrances, is capable of grounding knowledge in reality in a serious way. An adequate methodology of this kind would require a substantive, veridical, and explanatory justification of our claims to knowledge without imposing unnecessary requirements on the justificatory process. Borrowing from the title of Shapiro (1991), I will call the idea of such a methodology "*foundation without foundationalism*".³

I. Foundation without foundationalism

How shall we go about constructing a foundational, yet non-foundationalist, epistemic methodology? Our analysis in Section 2.1 provides a clue for answering this question. It suggests that the key to a viable foundational methodology is *freeing ourselves from the strict ordering requirement*, i.e., allowing non-strictly ordered grounding procedures. If we demand that the grounding relation satisfy the strict ordering requirement of the traditional methodology, the only route open to us is the foundationalist route or something like it. But why should the grounding relation be required to have this specific formal structure? Why should the rigid image of a foundation underlying past epistemologies restrict the search for a foundational

³ Shapiro uses "*foundations* without foundationalism" (plural), but I prefer "*foundation*" (singular) because the former has the connotation of "foundational studies", a notion that does not involve commitment to an actual foundation and is often used in contexts restricted to logic and mathematics.

methodology today? It is not hard to understand why philosophers would initially associate the foundational project with the idea of a strict ordering. Some of our most powerful grounding procedures, for example, logical proof, exemplify just such an ordering.⁴ But logical proof is just one tool for grounding knowledge. Why should the entire grounding process be modeled on just one tool, even as potent a tool as logical proof? The fact that some powerful grounding procedures have a certain formal feature does not mean that this feature is either necessary or sufficient for all grounding. Indeed, the severe problems plaguing foundationalism suggest that it is insufficient.

Relaxing the strict-ordering requirement, however, does not by itself ensure a better methodology. A prime example is *coherentism*, whose approach to the grounding relation is so relaxed as to give up the grounding-in-reality requirement, and with it the foundational project (as it is conceived here). Contemporary coherentists—Rescher (1973), Bonjour (1985), Williams (2001), and others—have brought coherentism closer to the foundational project by allowing some grounding of knowledge in reality; but their coherentism is still not foundational in our sense, since it does not *require* such a grounding, let alone require it for *all* branches of knowledge.

Foundationalism and coherentism, indeed, mark two extremes of the foundation–no-foundation divide. We can characterize their difference by saying that foundationalism is committed both to:

- (1) *the strict ordering of the grounding relation*
- and to:
- (2) *the groundedness of knowledge in reality,*

while coherentism is committed to neither. It is readily seen, however, that foundationalism and coherentism do not exhaust the array of possible positions with respect to (1) and (2). Foundation-without-foundationalism differs from both: unlike foundationalism it is not committed to (1), and unlike coherentism it is committed to (2). We can display the differences and similarities between the three approaches as follows:

	<u>Strict Ordering of the Grounding Relation</u>	<u>Grounding our System of Knowledge in Reality</u>
Foundationalism	required	required
Coherentism	not required	not required
Foundation w/o Foundationalism	not required	required

Figure 2.1

⁴ The procedure of logical proof allows rules of replacement which are symmetric. However, as noted earlier, these can always be eliminated, and the core structure of a logical proof is that of a strict, foundationalist-like ordering.

The key idea of the foundation-without-foundationalism approach is that there is no inherent connection between grounding our system of knowledge in reality and doing so in a strictly ordered manner. Our system of knowledge is connected to reality by a multifaceted relation (or network of relations), strictly ordered in some sections, not strictly ordered in others. All branches of knowledge, qua branches of *knowledge*, must be grounded in reality, but their grounding need not follow a single, strict, and rigid pattern. It might be argued that strict ordering is needed to avoid circularity and infinite regress. But it is widely acknowledged that not all circularity and regress are vicious. Indeed, it is not clear that philosophers' recoil from all forms of circularity and regress is not the outcome, rather than the reason for, the habitual strict-ordering requirement. Foundation-without-foundationalism is committed to avoiding *vicious* circularity/regress, but since not all non-strictly ordered grounding is viciously circular/regressive, the road is open to a non-foundationalist yet foundational methodology. (We will discuss circularity in greater detail later in this chapter.)

To sum up, foundation-without-foundationalism shares foundationalism's commitment to a foundation for knowledge while renouncing its rigid methods. It says that in grounding our system of knowledge we need not encumber ourselves with unreasonable restrictions and it grants us maximum freedom in designing and carrying out the grounding project. Unlike foundationalism, it does not determine in advance either the formal structure of, or the resources used in, each stage of the grounding process; and unlike coherentism, it does not give up on, or in any way compromise, the application of robust veridicality standards to our theories. Indeed, in one sense, at least, foundation-without-foundationalism, as it is conceived here, has a stronger commitment to veridicality than foundationalism. While foundationalism allows that some branches of knowledge (e.g., logic) are grounded exclusively in the mind, foundation-without-foundationalism, as we conceive of it here, requires all branches of knowledge to be significantly grounded in reality as well. Providing a sound grounding for our system of knowledge, however, is not making our system perfect or infallible. The grounding we are seeking is a solid grounding, yet our goal is realistic: we are seeking a grounding for human knowledge while recognizing both the inherent fallibility of such knowledge and our own fallibility. A methodology that makes such a grounding possible is valuable in the same way that a sound medical procedure is valuable. Such a procedure cannot provide an absolute guarantee of perfect health from here on, yet it has a considerable value for what it sets out to do: improve life.

In seeking to develop the new, non-foundationalist yet foundational methodology we may proceed in several directions. We may start, like Alston (1976), with foundationalism and weaken its requirements. We may start, like Bonjour (1985), with coherentism and strengthen its demands. We may start, like Haack (1993), with both foundationalism and coherentism, and build a bridge between them. And we may follow some other patterns starting with (or involving) foundationalism and/or coherentism, like Sosa's (1991) or Elgin's (1996). Very important advances in the foundational project have been made by all the above-mentioned authors (and

others), but one reason we have to go beyond the existent accounts is that they are largely confined to the foundation of empirical knowledge (i.e., to natural and social science), whereas our conception requires a foundational account that applies to all fields of knowledge—logic, mathematics, and philosophy included. Similarly, important attempts to provide a foundation for mathematics in the spirit of foundation-without-foundationalism were made by various philosophers, for example, structuralists such as Resnik (1981, 1982, 1997), Hellman (1989), Parsons (1990, 2008), Shapiro (1997), and others. But these, too, are restricted to a specific field. For that reason I have decided to pursue the project in my own way, based on the general principles of epistemic friction and freedom outlined in Chapter 1, which are applicable to *all* fields of knowledge.⁵

To make the project more manageable, I will focus primarily on theoretical knowledge. I am aware, of course, of the current tendency to start with non- or pre-theoretical knowledge (practical knowledge, common-sense knowledge, animal knowledge, etc.), and I share the view that a full understanding of knowledge must include both theoretical and non-theoretical knowledge. But while something important is gained by focusing on non-theoretical knowledge, something is also lost by neglecting theoretical knowledge. Given the aim of the present essay, I prefer to focus on the latter. A successful foundation for theoretical knowledge should be expandable to all knowledge, but this is something I will not attempt here.⁶

II. *Foundational holism*

In searching for an appropriate foundational methodology I will focus on two principles: the first is full commitment to the universal friction requirement; the second is flexibility in choice of steps for satisfying this requirement. More specifically, these principles say:

- (a) Every branch of knowledge, qua branch of knowledge, requires a substantial grounding in reality, i.e., justification of its veridicality.
- (b) The grounding relation need not be strictly ordered.

The first principle requires a *robust foundational* methodology: every discipline requires a substantial grounding in reality. The second principle licenses a *holistic* methodology: by renouncing the strict ordering requirement we allow the grounding relation to encompass a rich and diverse network of connections, both among units of

⁵ This leaves room, and indeed calls for, engagement with the views developed by the aforementioned authors. Doing this in the present essay, however, would enlarge it considerably and interfere with the presentation of a clear line of thought.

⁶ A note on internalism and externalism: since I am interested in a theoretical justification of knowledge, my project is largely internalist. But since justification, as I understand it, always involves factual correctness (correctness in the world), the project is externalist as well. Indeed, I think theoretical justification can be combined with other types of justification, including reliabilist justification (which I will not discuss here) and pragmatic justification (which I will).

knowledge and between units of knowledge and reality. Together, the two principles characterize a methodology that is both foundational and holistic. I will call this methodology “*foundational holism*”. Foundational holism affirms the grounding-in-reality requirement while allowing a rich holistic network to fulfill this requirement. It shares foundationalism’s conception of a foundation as providing informative explanation and veridical justification, but it allows us greater freedom in designing and carrying out the grounding project.

The question arises, however, whether a combination of holism and foundational epistemology is possible. Is holism compatible with the foundational project? More basically still, is holism itself an acceptable methodology? Let us start with the second question.

Is Holism an Acceptable Methodology? Some critics give a negative answer to this question. Holism, they say, is a flawed methodology. An especially compelling criticism of this kind is made by Dummett and Glymour who object to a version of holism which I will call “one unit” holism and which they identify (rightly or wrongly) with Quine. A central claim of this holism is that *the smallest unit of epistemic significance is our system of knowledge as a whole*, and, taking it to its “logical” conclusion, this claim means that there is *at most one* significant unit of knowledge. This, Dummett and Glymour rightly point out, is absurd:

[I]f a total theory is represented as indecomposable into significant parts, then we cannot derive its significance from its internal structure, since it has none; and we have nothing else from which we may derive it (Dummett 1973/81: 600⁷).

No working scientist acts as though the entire sweep of scientific theory faces the tribunal of experience as a single, undifferentiated whole (Glymour 1980: 3).

The point is that knowledge, on the one-unit approach, is an “all-or-nothing” affair: either we construct our system of knowledge all at once, or we do not construct a system of knowledge at all; either we test it in one fell swoop, or we do not test it at all; either we replace it in its entirety, or we do not change it at all; either we grasp (learn, communicate) it “in a single spasm of seamless cognition” (Fodor and Lepore 1992: 9), or we do not grasp it at all. We may liken our system of knowledge under one-unit holism to a huge atom or blob: no inner structure, no differentiation, no interrelations, nothing to work with. The very idea of a *system* of knowledge becomes meaningless. Truth, hypothesis, evidence, and inference lose their significance. Experiment is impossible. The connection between knowledge and rationality is undermined. *System of knowledge, truth, hypothesis*, etc., all require significant units of knowledge smaller than our system as a whole; but one-unit holism rules out the existence of such units.⁸

⁷ See also the longer discussion by Dummett (1973/81: 593–609).

⁸ For a further criticism of one-unit holism see Hofstadter (1954: 408–10). (Hofstadter also attributes this kind of holism to Quine.)

One-unit holism is indeed a barren holism. But *foundational holism* is *not*, and does not descend from, this type of holism. Foundational holism descends from *relational* or *network* holism, i.e., a version of holism that emphasizes the *interconnectedness* of knowledge (rather than its *oneness*). Interconnections require multiple units or modules of knowledge, and therefore relational holism is incompatible with one-unit holism. Relational or network holism is also commonly traced to Quine, but while one-unit holism is traced to Quine's claim that⁹ "[t]he [hence smallest] unit of empirical significance is the whole of science" (Quine 1951a: 42), relational holism is associated with his conception of our system of knowledge as an interconnected web.¹⁰

Foundational holism, however, requires more than just relations. The mere affirmation of interconnections allows loose, haphazard, and otherwise epistemically insignificant connections. But foundational holism, as I conceive it here, requires well-structured and significant connections. Foundational holism, therefore, is a *structural* rather than merely a relational holism. It views our system of knowledge as a highly and richly structured system, with multiple constituents standing in multiple epistemically significant relationships both with each other and with reality. This rich structure provides multiple avenues for obtaining, examining, and justifying knowledge (knowledge claims), making a non-foundationalist foundation for knowledge possible. Since structural holism stands in even greater contrast to one-unit holism than does relational holism, foundational holism, as a *structural* holism, is certainly exempt from Dummett's and Glymour's criticisms.¹¹ The question still remains, however, whether holism—even structural holism—is compatible with the foundational project.

Is Holism Compatible with the Foundational Project? Holism is sometimes thought of as a coherentist methodology, and as such as incompatible with the foundational project as it is understood here, especially with its universal grounding-in-reality requirement. I believe this view is mistaken. Although holism is compatible with coherentism, it is also compatible with a foundational project like ours, one that upholds the grounding-in-reality requirement and makes it universal. To see this, consider one of holism's paradigmatic metaphors—*Neurath's boat*.¹²

⁹ Throughout the work, I use **boldface** within citations to indicate my own emphases. (See fn. 3, Chapter 1.)

¹⁰ Some commentators distinguish between Quine's "meaning" and "confirmation" holism, regarding him as (in our terminology) a one-unit holist with respect to language but a network holist with respect to knowledge (see Fodor and Lepore 1992: Chapters 1 and 2). This, however, stands in some tension with Quine's view of language and theory (knowledge) as inseparable, which presumably implies that the two are governed by the same kind(s) of holism. Be that as it may, my concern here is to make clear that foundational holism is not subject to Dummett's, Glymour's, and Hofstadter's criticisms, whether they are right with respect to Quine or not.

¹¹ Structural holism bears some similarities to Dummett's molecularism, but it is a broader position, encompassing more connections than propounded by the latter, and not bound by all its restrictions.

¹² The metaphor is commonly traced to Otto Neurath in such passages as: "We are like sailors who have to rebuild their ship on the open sea, without ever being able to dismantle it in dry-dock and reconstruct it from the best components" (Neurath 1932: 92). It was later adopted by many philosophers, most notably

The Neurath boat metaphor represents our system of knowledge as a boat in the sea, and a problem with the system as a hole in the boat. A central difference between foundationalism and holism is that the foundationalist boat has a safe haven to return to while the holist boat does not. To fix a hole, the foundationalists bring their boat back to shore, make the needed repairs in the safety of their docks, and only then launch it back to sea. But the Neurathian sailors remain on their boat and, using whatever resources they have then and there, patch the hole and continue sailing. Once they have mended one area of the boat, they use it as a temporary foothold for inspecting other areas and repairing other holes. At some point in this process they revisit the original hole, examine the patch, and utilizing new resources accrued or created on the boat since the patch was installed, strengthen it or replace it by a sturdier patch. In this way any section of the boat can in principle be repaired, strengthened, and improved, and any section can serve as a temporary foothold.

The Neurath boat metaphor is often interpreted as a coherentist metaphor (see e.g., Sosa 1980a)¹³. But it is important to distinguish the metaphor's anti-foundationalist features from its coherentist features (on such interpretations).

A. Anti-foundationalist features:

1. There is no Archimedean standpoint and no need for one.
2. There is no one fixed procedure or pattern for repairing holes in the boat.
3. The repair process is in certain ways circular, but this circularity is constructive rather than destructive (it enables us to strengthen existent patches or replace them by better ones).¹⁴

B. Coherentist features:

1. The boat is separated from the world.
2. The boat is self-sufficient.

While it is possible to construe Neurath's boat as having both the (A) and (B) features, and thus as representing a coherentist anti-foundationalism, it is also possible to construe it as having only the (A) features, and thus as compatible with a strong foundational project.

Indeed, the boat image itself suggests a foundational interpretation (in our sense). In the first place, the boat does not hover in empty space, but lies in real, external, resistant waters; and its sailors must take the nature of this external medium into account (they cannot patch a hole in the boat with a material that is not water-resistant).

Quine, and received a life of its own, independent of Neurath's original intentions. Here I consider the Neurath-boat metaphor as standing on its own and as open to diverse interpretations.

¹³ It should be noted that Sosa's conception of a foundation is both broader and weaker than ours, and this enables him to regard coherentism as a (weak) foundational methodology.

¹⁴ More on this shortly.

In the second place, Neurath's boat is not a dinghy afloat in a still pond, but a ship in the sea, exposed to the pounding of waves, the buffeting of winds, the flooding of water, etc. So its sailors have to take into account not just one feature of external reality but many. Moreover, Neurath's boat is not a pleasure boat or a boat aimlessly traversing the seas. It is a vessel on an expedition, a *Beagle* or a *Victoria*, whose task is to study the external sea and its environs. All these suggest that Neurath's boat is not isolated or separated from the external world, but on the contrary, is oriented toward it and affected by it. Furthermore, being on a prolonged mission, and one started at sea rather than at a richly supplied docking facility, the boat is not self-sufficient. Where would its sailors get their nutrition, fuel, raw materials, and other resources from, except from the water, the wind, the rain, the fish, the debris, and other elements of their external environment? Neurath's boat is thus invested in the world, directed toward the world, and constrained by the world. Its mission is to study the world, and its strategy is to use any available resources, internal or external, in any possible way, strictly ordered or not. As such, Neurath's boat is not only compatible with *foundational holism*, but also representative of it.

Foundational holism, as represented by the Neurath boat metaphor, incorporates the two main principles of the foundation-without-foundationalism methodology: (i) knowledge—all knowledge—requires a grounding in reality, and (ii) the grounding of our system of knowledge in reality relies on a rich network of grounding relations, not all strictly ordered. Using Sosa's (1985) terminology of "intellectual virtues", we may say that foundational holism allows us to use the full range of intellectual virtues in the grounding process. The rich network of interconnections provides multiple opportunities for making cognitive contacts with the world, opportunities for the world to make itself known to us and for us to ground our knowledge in it.

What distinguishes foundational holism from other forms of holism is its putting the holistic network in the *service* of the foundational project. We may describe the distinctive features of foundational holism as follows:

- (a) Foundational holism involves, in addition to the *traditional* holistic idea of a rich network of connections *among units of knowledge*, the *new* idea of a rich network of connections *between units of knowledge and reality*.
- (b) Foundational holism not only *permits* connections between our system of knowledge and reality (as some reformed versions of coherentism do), it makes such connections *mandatory*. And it makes them mandatory not just for some, but for *all* branches of knowledge.
- (c) Although foundational holism regards knowledge as grounded not just in reality but also in the mind, it never regards the grounding of knowledge in the mind as a substitute for its grounding in reality. All knowledge is grounded both in reality and in the mind, and the former is *integral* to its *veridicality*.

These three features separate foundational holism from coherentist holism. While coherentist holism, and indeed most traditional forms of holism, focus on the

existence of an intricate network of connections among diverse cognitive units (concepts, terms, statements, theories, disciplines, etc.), foundational holism emphasizes the existence of an additional network of connections as well. The acquisition and grounding of knowledge involve, in principle, not just a large and intricate network of connections among diverse cognitive units, but also a large and intricate network of *cognitive routes from mind to reality*—a large and open-ended array of methods used for cognitively reaching reality, constructing theories of various aspects of reality, and grounding these theories in reality. In this sense foundational holism is a *double holism*, characterized by the existence of *two* rich networks of connections, indeed, two *interconnected* networks.

Furthermore, foundational holism is not a piecemeal methodology. Holism, as captured by the Neurath boat metaphor, is sometimes associated with anti-methodism and piecemeal philosophy. But whether this was Neurath's original intent or not, it is not inherent in the metaphor itself. A back-and-forth process of investigation can be highly systematic, and while the metaphor itself is compatible with piecemeal philosophy, it is not restricted to such philosophy. The present approach, motivated by (among other things) the friction requirement of substantiveness, which includes systematicity, is strongly committed to the latter.

The grounding project itself is a joint project of philosophy and other disciplines. Philosophy's job includes constructing a general methodology for grounding knowledge both in reality and in the mind and a general model of such a grounding, explaining the foundational differences between different branches/types of knowledge, and providing an actual grounding for distinctly philosophical (or partially philosophical) fields of knowledge (e.g., logic). The grounding project as a whole, however, is an interconnected project, and its philosophical chapter is neither isolated from, nor has a greater authority than, those of other sciences. As a result, philosophy's standing in the grounding project is radically different from its traditional standing as a "first philosophy". At the same time, its standing is also different from what radical naturalists take it to be. Philosophy is not transmuted into a branch of empirical science, be it psychology, neurology, sociology, or any other science; nor is its role limited to giving a helping hand to these sciences.

One possible challenge to foundational holism, formulated in terms of the Neurath boat metaphor, is: How do you know that there is no part of the boat which is so essential to it that, if it breaks down, the boat will be completely destroyed and there will be no place to stand and fix it? The answer is that we do not know this in advance. But the foundationalist does not know in advance that there is such a part either. Nor would it help him if he did. His methodology, as we have seen above, cannot ground our system of knowledge, whether or not it has such crucial parts. The possibility of a solid foundation is thus not guaranteed in advance regardless of what methodology is used, but the development of an effective yet flexible foundational methodology increases the likelihood of such a foundation.

Our next task is to delineate a model of knowledge exemplifying the foundational-holistic methodology. This model will be foundational, holistic, and structural. It will represent the grounding of our entire system of knowledge both in reality and in the mind, set high standards of truth and substantiveness for all disciplines, and enable us to see how “resistant” disciplines (such as logic and philosophy) can satisfy these standards. Before I turn to this task, however, let me consider, once again, the question of circularity.

Circularity: Destructive, Trivializing, Constructive. Foundational holism frees us from the quest for an infallible, absolute, once-and-for-all foundation, yet it does so without renouncing the goal of a substantive, explanatory theoretical foundation. A central step toward the achievement of this goal is the renunciation of the traditional ordering requirement on the grounding relation, which opens the door to holistic grounding patterns (alongside the traditional pattern). One result of this move, however, is license to include a certain measure of circularity in the grounding process. But should any measure of circularity be permitted? Now, there is no question that circularity can be destructive, giving rise to paradoxes such as Russell’s or the Liar. Yet it is just as clear that circularity is often innocent. After all, our paradigm of circular reasoning, “P; therefore P”, is logically valid, and as such does not, and cannot, produce error.

Innocent circularity, nevertheless, can be trivializing. Leaving no room for friction, it can render our theories empty, make our answers question-begging, and undermine our claim to a genuine grounding or justification. In recent decades philosophers have focused on a particular kind of circularity relevant to the foundational project: *epistemic circularity* (see e.g., Alston 1986, Sosa 1994, 1997, 2009, and Lammenranta 1996, 2009). Epistemic circularity arises when one tries to establish, justify, or provide a foundation for X in a way that involves, either explicitly or implicitly, reliance on (assumption, presupposition of) X.¹⁵ One significant problem raised by epistemic circularity is what Cohen (2002) calls the problem of “easy knowledge”. The problem is, essentially, that using knowledge of type X to provide a foundation for field of knowledge X makes the foundational task “too easy”. In our terminology, the problem is that allowing epistemic circularity trivializes our foundational results or renders the foundational process *frictionless* (lacking adequate friction). It is not clear, however, that all innocent circularity is frictionless or otherwise problematic. For example, the claim that circularity *precludes the discovery of error* appears to be exaggerated. One might think that we cannot hope to discover errors in our logical principles given that we have to use these very principles in the discovery process. But using them is one thing, and being committed to them, or being blinded by our use of them, is another. Take, for instance, Russell’s discovery of

¹⁵ Alston’s paradigmatic example of epistemic circularity is the establishment of sensory perception as a reliable source of knowledge by means of an argument whose premises presuppose the epistemic reliability of such perception. (See his 1986: 8–10.)

a paradox in Frege's logic. In discovering "Russell's Paradox" Russell had to use some logic. Which logic did Russell use? Clearly, he had to use a quite powerful logic, a logic more powerful than, say, sentential logic. But he could not have used his own type-theoretic logic, which he developed *after* discovering the paradox and largely as a means of avoiding it. Nor could he have used standard first-order logic plus Zermelo's axiomatic set theory. In all likelihood, Russell used something on the order of Frege's logic, yet this did not prevent him from discovering the paradox. How? It is reasonable to surmise that whatever logic Russell used, he used it flexibly, dynamically, critically, and intelligently—holding off some parts, switching from part to part, and so on—so the paradox could come to light. A similar, perhaps more perspicuous case, is that of Grelling's discovery of the "heterological" paradox, which was in all likelihood made in a language susceptible to this paradox. Likewise, the discovery of the Liar paradox, and many of its proposed solutions, were made using languages that are not immune to semantic paradoxes. So the use of methods involving circularity does not always prevent the discovery of error.

Indeed, it is reasonable to surmise that, rightly used, circularity—like other non-traditional patterns—may *increase our cognitive powers* in a variety of ways. In this respect methods involving some measure of circularity are similar to methods involving an element of reflexivity. From Cantor's and others' use of the diagonal method, to Gödel's use of syntax to represent syntax, and Henkin's use of syntax to build a semantic model for syntax, the use of reflexive methods has been associated with ingenious achievements in fields such as metalogic and mathematics. In fact, the very idea of developing a metalogic, ML, for a logic, L, is predicated on the premise that we can use tools developed by mathematical theories formulated within the framework of L to draw results about L itself, including results (like Gödel's) whose impact is felt far beyond the specific system L.

I will call the circularity involved in furthering our epistemic goals "*constructive circularity*". Constructive circularity is central to logic and philosophy in a variety of fields. Rawls's method of reflective equilibrium in the theory of justice, Glymour's bootstrap method in the philosophy of science, Gupta and Belnap's revision method in the theory of truth, all involve constructive circularity. Rawls's (1971) method, modeled after a method suggested by Goodman (1953/5), emphasizes *back and forth* movement between particular judgments and general principles, leading to an *equilibrium*. Goodman, who introduced this method as a method for justifying logic, is explicit about its circular nature. Describing the process, he says:

This looks flagrantly circular. I have said that deductive inferences are justified by their conformity to valid general rules, and that general rules are justified by their conformity to [particular] valid inferences. But this circle is a virtuous one. The point is that rules and particular inferences alike are justified by being brought into agreement with each other. A *rule*

is amended if it yields an inference we are unwilling to accept; an inference is rejected if it violates a rule we are unwilling to amend. The process of justification is the delicate one of making mutual adjustments between rules and accepted inferences; and in the agreement achieved lies the only justification needed for either (Goodman 1953/5: 64).¹⁶

Glymour's (1980) bootstrap method licenses the use of a scientific hypothesis to aid in the confirmation of a theory it belongs to, or of some parts of this theory. And Gupta and Belnap's revision method (1993) provides a way of making sense of, and dealing with, circular concepts by identifying a fruitful form of definition associated with these concepts.

In the present context, we can point to two effective devices for using circularity constructively, avoiding trivialization. One important device can be described as *bringing new elements into the (circular) mix*. Thus, to provide a foundation for X, we may use some elements of X in combination with other elements, external to X, and possibly new combinations of old elements involving or related to X. In Rawls's theory of justice, this leads to a *wide* reflective equilibrium, where background and newly discovered knowledge is added as input to the equilibrium-generating process.¹⁷ It is significant that Rawls conceives of wide reflective equilibrium as a *non-conservative* method (Rawls 1974–5: 7–8), i.e., a method leading to results that are essentially richer than the general principles and particular judgments used at the outset. This is a central feature of *constructive* circularity in our sense. In Glymour's case, the introduction of *new elements* into the circular mix results in a *non-pure* bootstrap process. A particular case of non-pure bootstrap involves two theories of (roughly) the same level of generality, developed (grounded) in tandem. We use resources from Theory A to further develop (ground) Theory B, then we use some of the new elements of Theory B as resources for further developing (grounding) Theory A, and so on. We may call this process "*double bootstrap*". In Chapter 10, Section 8 we will see an example of a double bootstrap, involving logic and mathematics.

Another potent device for avoiding the trivializing/undermining effects of circularity (infinite regress) is the *switching of perspectives*—switching standpoints on Neurath's boat, ascending to a higher level of discourse (a metalanguage or a metatheory), moving to another theory on the same level (a background theory), etc. A paradigmatic example of ascent to a higher level is Tarski's theory of truth. Tarski defined truth in stages. Starting with a so-called object language, OL, his recursive definition of truth for OL is formulated in a more powerful metalanguage, ML, whose own truth predicate belongs to its own metalanguage, MML, and so on. It may be objected that this move trivializes/undermines Tarski's definition due to

¹⁶ My own suggestion of how to arrive at a foundation for logic is different from Goodman's (as the reader will see in Chapter 10). But as an example of constructive circularity Goodman's suggestion holds.

¹⁷ See Rawls (1974–5) and elaboration in Daniels (1979).

something akin to infinite regress (see e.g., Rescher 2010)¹⁸, but in my view this objection is unwarranted. The dynamics of Tarski's definition, viewed from a holistic, Neurath boat perspective, may be described as follows: At the outset, the truth predicate of OL, T_{OL} , is an undefined, loosely construed, merely intuitive notion of ML, and the truth-predicate of ML, T_{ML} , is likewise an undefined notion of MML. At the next stage, we ascend to ML and use its powerful formal apparatus to provide a rigorous definition of T_{OL} . Next we ascend to MML and use its still more powerful formal apparatus to turn the intuitive notion of truth for ML, T_{ML} , into a rigorous notion, and so on. At some point we step aside from the "original" Tarskian hierarchy to some background theory from which we can view that hierarchy as a whole, and in observing it we understand how the imprecise use of truth in a given stage is made precise by ascent to the next stage, i.e., we understand the *general principle of truth* as it appears in the original hierarchy. Once we understand how truth works in the original hierarchy, we can generalize our understanding to an arbitrary Tarskian hierarchy, including the hierarchy we realize we are standing in when studying the original Tarskian hierarchy. At this point we attain a general understanding of truth based on Tarski's definition, in spite of the infinite hierarchy it is associated with.

Both devices are based on the principle that the development/foundation of knowledge is a *dynamic process*, requiring a continuous exercise of epistemic freedom. The result is that circularity is rendered *partial*: in developing/grounding a field of knowledge X we may use elements of X essentially, but at each stage we use only part of X, at different stages we use different parts of X, we use other, new elements in addition to X, and we are always open to the possibility of revising X. Similarly, by changing our point of view we have multiple perspectives on a given circularity, minimizing its impact. This dynamic approach contrasts with the static approach characteristic of most discussions of epistemic circularity in the literature. The circularity problem is usually presented as arising in a static situation involving a completed foundation X (or a proposal for such a foundation), rather than in a holistic, dynamic situation involving multiple connections between knowledge and reality, as suggested here. Others, too, have recently suggested something along the latter lines as a key to mitigating the epistemic circularity problem. Thus,

¹⁸ Focusing on the "only if" part of the T-biconditional, Rescher formulates this objection as follows:

Consider the propositional regress

$$\text{"p" is true} \leftarrow p$$

This is the truth regress or Tarski regress... [One possible way to view it is] as laying down preconditions for the truth-status determination of regress elements. You can only determine the truth status of a member if you have already pre-determined the truth status of its predecessor, with that predecessor itself being in the same boat. Here, then, we have a Catch-22 situation of unsatisfiable preconditions, and the regress is consequently vicious and paradoxical (Rescher 2010: 30).

speaking about “reflective” justification—his correlate of our “theoretical” justification—Sosa says:

The right model for understanding reflective justification is *not* the linear model whereby justification is a sort of liquid that flows through some pipe or channel of reasoning, from premises to conclusion. (Such flow is linear, unidirectional; the pipe or channel “transmits” the justification—or warrant, or epistemic status.) A better model is rather that of the **web of belief**, whereby the web is **properly attached to the environment**, while its nodes can also gain status through **mutual support**....By basing beliefs on other beliefs the rational weaver weaves a web each member of which is held in place *in part*...through its being based on certain others, directly or indirectly. There is no apparent reason why such basing should be regarded as either causally or normatively asymmetrical, no reason why many beliefs could not constitute webs in which *each* mode is in place by being based *partly* on the *others*. What is more, each might thus gain its epistemic status partly through such relations with the others, where the whole web is also **attached to the world** through the causal mechanisms of perception and memory.

...Of course, when in psychology or neurobiology or cognitive science or common sense we modify an epistemic commitment, whether implicit or explicit, we do so based on beliefs that we acquire through our relevant commitments already in place, prominently those involved in our perceptual acquisition of information. There is hence an **inevitable circle** involved in the way we come to modify and hold our relevant perceptual commitments, whether implicit or explicit. For we hold them, and sustain them over time, based on **continuing** observations, and on particular perceptual beliefs, which are themselves based on the now installed, and **perhaps modified** commitments. There seems no special viciousness pertaining to the nodes of our web constituted by these commitments (Sosa 2009: 239–40).

And earlier he emphasizes that “*connection with the relevant externalia*” is crucial for the success of this process (Sosa 2009: 76).¹⁹

Returning to our own foundational holism, it is important to note that while we accept some forms of circularity as harmless and even as advantageous, we do not give circularity (regress) a blanket endorsement. On the contrary: circularity, from our perspective, has both the potential of being destructive and the potential of being constructive; as such, it must be handled with care. But circularity is too powerful to be given up altogether. The epistemologist must be constantly on guard to avoid destructive uses of circularity, yet she must not neglect her search for new, constructive uses of circularity. Reid compared circular justification to “[t]rying to settle whether a man is honest or not by asking him” (Reid 1785: 259). But if we use multiple other sources to establish a man’s honesty, then including his own testimony is, in principle, legitimate, and can be helpful, for example, by directing us to additional external sources.

¹⁹ For a dynamic conception of the wide-reflective-equilibrium method see e.g., Daniels (1979: 267).

III. *A note on grounding*

In this essay I use “grounding” in a broad sense, synonymous with “foundation”, so that providing a grounding for a claim, a theory, a discipline, or a system of knowledge X involves developing a substantive, explanatory, largely normative, philosophical account of the fundamental features and veridicality of X, focusing on its dual dependence on the world on the one hand and on the mind on the other. From this perspective, we can view the seven principles of epistemic friction delineated in Chapter 1, Section 2 as normative conditions on an adequate grounding of knowledge.

In the current literature it is customary to distinguish between (i) epistemic grounding, (ii) metaphysical grounding, and (iii) methodological grounding. Take logic, for example; (i) involves answering such questions as: “What justifies us in accepting a given logical theory or in claiming that one sentence follows logically from another?”, (ii) involves answering such questions as: “What facet of reality are logical consequences grounded in? How does being grounded in this facet endow logic with the strong modal force it is required to have?”, and (iii) involves answering such questions as: “Given the basicness of logic, how can we ground it in the sense of (i) and (ii)?” From our perspective, these questions are closely connected. For example, to ground logic epistemically we need to show that it is grounded in a facet of reality that endows it with certain requisite features (e.g., strong modal force), something that involves a metaphysical grounding. The present chapter has developed a new methodology of grounding, foundational holism, which is intended to replace the traditional foundationalist methodology.²⁰

In Part II, I will use the foundational-holistic methodology to delineate a general model of knowledge based on the principles of epistemic friction and freedom. A distinctive characteristic of this model is its highly dynamic structure and the wide array of connections between theory and the world it countenances. This dynamic structure makes it possible for *all* branches of knowledge to be grounded both in the world and in the mind—logic in certain structural features of the world, and observational science in a conceptual system generated by the mind. In Part III, I will develop a theory of truth for this dynamic model and explain how it navigates

²⁰ The recent literature on grounding—e.g., Fine 1995 and 2012; Schaffer 2003 and 2010; Rosen 2010; and Sider 2011—is largely devoted to metaphysical grounding. Like Rosen I believe that speaking in terms of grounding is philosophically fruitful. Like Schaffer (2003) I am suspicious of the assumption that there is a single, absolute, fundamental level of reality (this assumption is uncomfortably close to some of the problematic tenets of the foundationalist outlook). Like Fine (2012) I associate grounding with explanation and I see it as coming in a variety of strengths and “flavors” (for example, metaphysical grounds are modally stronger than physical grounds, and grounding can be either factive or non-factive—in the present terminology, grounding in reality is factive, grounding in the mind is not). And like Sider I believe that metaphysical questions are substantive and that a good answer to such questions “carves the world at its joints”.

the universal friction requirements of substantiveness and grounding-in-reality. This theory develops a non-traditional conception of correspondence, one that extends the correspondence standard to abstract fields of knowledge (such as mathematics) without giving rise to either the problems of empiricism or to those of Platonism. In Part IV, I will further use the foundational-holistic methodology to delineate an outline of a foundation for logic. This foundation grounds logic both in the world and in the mind, explains why, in order to perform its designated function, logic must take the world into account, and identifies a facet of reality which provides a grounding of the requisite kind for logic. The interplay between logic and mathematics in that foundation is, as I have noted earlier, an example of constructive circularity.

PART II

A Dynamic Model of Knowledge

3

Quine's Model of Knowledge

An Inner Tension

3.1 The Initial Promise of Quine's Model

Having set the ground for the construction of a general model of knowledge based on the principles of epistemic friction and freedom, we are ready to put our holistic methodology to use. In certain respects, we are seeking a model that sets higher standards for our system of knowledge than past models. For example, the requisite model will require that not just the natural and social sciences, but also mathematics, logic, and philosophy, be grounded in reality, and that not just abstract knowledge, but also experiential knowledge, have a substantial conceptual grounding. All disciplines, in this model, will be subject both to high standards of truth, objectivity, and veridical justification, and to high standards of conceptualization, unity, and substantiveness. But can all disciplines satisfy all these standards? It is not easy to see how they can—how logic, for example, can satisfy high veridicality standards, or how philosophy can satisfy high substantiveness standards. This naturally tempts philosophers to weaken some of these standards for some disciplines, and possibly exempt them altogether. It is thus not surprising that almost all epistemologists exempt logic from the grounding-in-reality requirement, and that many release mathematics and philosophy from this requirement as well. Similarly, it is not surprising that some philosophers—quietists, deflationists, and minimalists—release philosophy (or some of its theories, or substantial parts of some of its theories) from the substantiveness requirement. But one should be wary of the temptation to treat these disciplines as incapable of assuming the normal responsibilities of a *field of knowledge*. In particular, we should be wary of yielding to this temptation prior to exploring new possibilities, both methodological and structural.

Our search for a new epistemic methodology in Chapter 2, Section 2 was a way of combating this temptation on the methodological level. In this Part, I will search for a model of knowledge that is so structured as to allow all fields of knowledge significant access both to the mind and to reality. A central feature of this model will be its rejection of the traditional division of units of knowledge into factual and non-factual. All fields of knowledge, qua fields of knowledge, are both significantly factual and significantly extra-factual.

Based on these (as well as earlier) considerations, the model we are looking for has the following building blocks:

1. A systematic, yet flexible structure.
2. A broad, significant interface with reality.
3. A broad, significant interface with the mind.
4. A rich, holistic network of interconnections, encompassing both connections among units of knowledge and connections between units of knowledge and reality.
5. Non-bifurcation of units of knowledge into units grounded in reality and units grounded in the mind (or units substantially grounded in reality and units exclusively grounded in the mind). All units of knowledge are substantially grounded in both.
6. A broad, yet substantial conception of mind, reality, unit of knowledge, truth, evidence, and justification.
7. High standards of substantiveness and veridicality for all disciplines.
8. A flexible and dynamic conception of epistemic inquiry.
9. A significant role for active freedom in the project of knowledge.

In searching for a model of knowledge with these building blocks, one model in particular attracts our attention. This is Quine's model of knowledge, developed in "Two Dogmas of Empiricism" (1951a) and other early writings, and partly incorporated in his later writings. Quine's model has a systematic structure (a *center-periphery structure*), a significant interface with reality (*periphery*), and a significant interface with the mind (*center*). It has a rich holistic network, connecting diverse units of knowledge to each other (and ensuring, among other things, a non-rigid structure). Moreover, the traditional bifurcation of units of knowledge into units grounded in reality and units grounded solely in the mind is rejected by his model (as part of an all-out rejection of central traditional dichotomies). The model admits both experiential and abstract units of knowledge, and its criterion of ontological commitment sanctions, in addition to observable physical objects, also theoretical scientific objects and mathematical objects. Finally, the model attributes a central role to pragmatic-conventional principles in the project of knowledge, thereby recognizing the centrality of (at least some aspects of) epistemic freedom to this project. All these make Quine's model a promising starting point for our endeavor.

But Quine's model has glaring shortcomings for this endeavor as well. To see these, let us begin with an inner tension in Quine's model.

3.2 An Inner Problem in Quine's Model

In his 1973 paper, "The Significance of Quine's Indeterminacy Thesis", Dummett makes a provocative claim: Quine rejected the analytic-synthetic distinction in "Two Dogmas of Empiricism" on the ground that no coherent and philosophically

significant content has ever been assigned to it. Yet, ironically, the distinction has come to life for the first time in the very account that was meant to replace it: Quine's holistic model of knowledge with its center-periphery structure. Dummett's claim is based on two observations: (a) Quine's center-periphery model provides us with notions in terms of which the analytic-synthetic distinction is well defined: "an analytic sentence is one such that no recalcitrant experience would lead us to withdraw our assignment to it of the value true, while a synthetic one is one such that any adequate revision prompted by certain recalcitrant experiences would involve our withdrawing an assignment to it of the value true" (Dummett 1973: 375). (b) This version of the distinction cannot be rejected as empty, since the claim that there are no synthetic statements and no analytic statements (in the above sense) conflicts with the center-periphery model. Dummett's conclusion: Quine's claim that the analytic-synthetic distinction is unfounded conflicts with his claim that our body of knowledge has a center-periphery structure.¹ I will not discuss Dummett's account of the conflict here: for one thing, my point of view is somewhat different from Dummett's; for another, I have some qualms about certain aspects of Dummett's argument. Instead, I will present the conflict in my own way and from my own perspective, namely, as a conflict between two epistemic methodologies: the one associated with the "no analytic-synthetic bifurcation" thesis (NAS), the other with the "center-periphery" thesis (CP).²

A Methodological Conflict. The most basic dichotomy underlying the analytic-synthetic distinction is that of statements grounded in *matters of fact* and statements grounded in something *other than fact*. What this other thing is varies from one philosophical theory to another, but Quine's attack applies equally to older as well as to newer renderings of the distinction. Methodologically, the most important feature of the analytic-synthetic thesis is the *bifurcation of our standards of knowledge*: the standards for the acceptance, justification, revision of synthetic items of knowledge are different from the standards for the acceptance, justification, revision of analytic items. Quine's revolution consists in rejecting this bifurcation. Every item of knowledge is subject to both standards: logic may be revised in response to experience; observation statements may be preserved based on pragmatic considerations.³

¹ Dummett also raises this issue in his (1973/81, Chapter 17).

² (i) By offering my own version of the conflict, I do not wish to diminish my debt to Dummett; needless to say, Dummett is not responsible for my version.

(ii) Below I use "NAS" and "CP" both for the theses and for the bifurcations (distinctions, divisions, dichotomies) themselves.

³ (i) I am not the first to view the philosophical weight of the analytic-synthetic distinction as epistemic. An earlier proponent of an epistemic reading of this distinction was Frege:

[The] distinction... between... synthetic and analytic... concern[s], as I see it, not the content of the judgment but the **justification for making the judgment**.... When a proposition is called... analytic... this is... a judgement about the ultimate ground upon which rests the **justification** for holding it to be true (Frege 1884: 3).

Now, CP breaks this uniform picture of knowledge by positing a structure with two essentially different, if interconnected, zones: an inner zone—*center*—based on one set of standards, and an outer zone—*periphery*—based on another. Elements in the center are manipulated using pragmatic standards, elements in the periphery—using evidential standards. Elements located in the periphery stand in a privileged relation to reality that elements located in the center are excluded from. Thus, granted that a conflict with experience can be resolved by changes anywhere in the system, there is still an essential difference between changes occurring in the center and changes occurring in the periphery. The ground for changing statements in the periphery is the need to square these statements themselves with experience; the ground for changing statements in the center is the need to square other statements, namely statements in the periphery, with experience. Truths located in the center are held solely in virtue of their contribution to the overall working of the system, hence due to pragmatic considerations. Truths located in the periphery are held partly because of their contribution to the system as a whole, but largely, and most importantly, due to their own direct link to experience, i.e., based on factual, objective considerations. Take logic, for example. There is no room in Quine's model for a conflict between a logical statement, ℓ , and an external event, e , analogous to, say, the conflict between the peripheral statement, "It rained in San Diego on Monday, March 1, 1999", and the event of sunny skies in San Diego on that day. Reality never impinges as strongly on logic as it does on experimental science, and logic, therefore, never occupies the periphery. We may change our logic in response to conflicts with experience, but only pragmatically and via a chain of connections starting and ending in the periphery. The bone of contention between CP and NAS is the role of fact and convention in knowledge: while a principled division of knowledge into conventional and factual constituents is denied by NAS, such a division is affirmed by CP. By superimposing a center-periphery structure on the unified domain of knowledge

(ii) Some of those who share my view that the main significance of NAS is epistemic have suggested that Quine's target should have been the apriori-aposteriori thesis, which is straightforwardly an *epistemic* thesis (unlike NAS which is literally, at least, a linguistic thesis). Although I do not deny that Quine rejected the apriori-aposteriori thesis, I think the analytic-synthetic thesis should have been his main target. This is because the analytic-synthetic bifurcation induces a deeper and more basic epistemic bifurcation than the apriori-aposteriori bifurcation. Let me explain: The analytic-synthetic bifurcation divides units of knowledge into units that are *factual* and units that are *not factual*, whereas the apriori-aposteriori bifurcation divides units of knowledge into units that are *empirically factual*, and units that are *not empirically factual*. But the category of "fact" is broader and more fundamental than that of "empirical fact". Thus, while the analytic-synthetic distinction creates a very fundamental cleavage—a cleavage between factual knowledge (in general) and other types of knowledge, the apriori-aposteriori distinction creates a less fundamental cleavage—a cleavage between a particular type of factual knowledge and other types of knowledge. As a result, NAS rejects a more fundamental bifurcation than the *No Apriori-Empirical Distinction* thesis. I should acknowledge, however, that Quine does use NAS more narrowly than indicated above, identifying "fact" with "experience", hence "empirical fact", and non-fact with "linguistic conventions or pragmatic considerations". Hence, for him, the two theses do coincide. But my own interest here is in the former (broadly construed) rather than in the latter.

postulated by NAS, Quine has *de facto* incorporated certain significant elements of the analytic-synthetic dichotomy in his epistemology.

3.3 Objections and Responses

Before turning to possible solutions, I would like to consider three likely objections to the reality of the conflict.

Objection 1. CP cannot be said to introduce an analytic-synthetic distinction into Quine's model, since such a distinction requires an absolute division of statements into two mutually disjoint groups but Quine's model represents our body of knowledge as a continuum. The center-periphery duality is gradual, while an analytic-synthetic dichotomy requires sharp boundaries.

Response. First, I would like to emphasize that on my construal, the center-periphery duality does not generate a full-scale analytic-synthetic distinction. At issue is a significant overlap of the two distinctions, not a complete identity (or even inclusion).

Second, the claim that the analytic-synthetic distinction does not allow a graduated progression may be true to its letter but is not true of its use, especially at the hands of Quine's main targets, the logical positivists. Explanation: traditionally, the analytic-synthetic distinction is a distinction between *whole* statements; but for the logical positivists, a more basic distinction exists between *constituents* of statements. This secondary distinction plays an especially important role in Carnap's reductionist project in the *Logical Structure of the World* (1928)⁴. Science, in Carnap's construction, is built out of basic elements of two types: purely experiential elements (in Carnap's construction, primitive "elementary experiences" and the primitive relation of "recollection of similarity"), and purely conventional elements (logico-mathematical operations and linguistic operations like definition). The construction proceeds in a series of stages: in stage one, purely conventional operators are applied to purely experiential elements, generating mixed elements of level one; in stage two, conventional operators (and mixed operators obtained in stage one, if any) are applied to elements of level one, generating mixed elements of level two, and so on. In this way we progress from elementary experiences to pairs of elementary experiences, equivalence classes of elementary experiences (under the relation of recollection of similarity relative to a given agent), and so forth. Somewhere along the way we obtain qualities (quality-structures), observable physical objects, theoretical objects, and so on. The result is a continuum of elements and, corresponding to it, a continuum of statements. If we order these statements according to their balance

⁴ Note: throughout this essay I will assume the prevalent interpretation of Carnap, which was arguably largely shaped by Quine's criticisms, rather than any of the newer interpretations (e.g., Friedman's 1999), since the former is better suited for explaining Quine. Whenever a conflict with some reader's interpretation arises, she may read "Carnap" as abbreviating something like "Quine's Carnap".

of experiential to conventional constituents (based on one criterion or another), we arrive at a progression of statements stretching from the purely synthetic to the purely analytic. We could, of course, draw a line (mark a region) somewhere along the continuum: statements on one side of the line (inside the designated region) would be “absolutely” analytic, all the rest “absolutely” synthetic; but such a division would still, at bottom, be graduated.

Thus, consider Carnap’s analysis of simple synthetic statements, say, “All dogs are carnivores”. This statement, on Carnap’s analysis, is doubly mixed: on the one hand, it contains both analytic and synthetic constituents (the analytic operators “all” and “if... then”, and the synthetic predicates “dog” and “carnivore”); on the other hand, its synthetic constituents themselves are decomposed into analytic and synthetic subconstituents. The concept “dog”, for example, is constructed by successive applications of analytic operations, first to purely synthetic elements⁵ and then to mixed elements, generating a gradual progression of concepts (and, correspondingly, objects) from the purely synthetic to the increasingly analytic. Speaking objectually, this progression includes:

- (1) Elementary experiences and the relation of recollection of similarity between them (primitive).
- (2) *Dog* sense-qualities: for example, the color brown. (Equivalence classes of elementary experiences under the relation of recollection of similarity.)
- (3) *Dog* sense-qualities-at-space-time-locations: for example, brown in the kennel at *t*. (*n*-tuples whose first four elements are real numbers—representing space-time locations—followed by one or more *dog* sense-qualities.)
- (4) *Dog* points-of-the-perceptual-world. (Classes of *dog* sense-qualities-at-a-location.)
- (5) *Dog* states. (Classes of *dog* perceptual-points.)
- (6) Individual *dogs*. (Classes of dog-states.)
- (7) The general object or species *Dog*. (The class of all individual *dogs*).⁶

This series marks a *graduated progression* from purely synthetic elements to elements whose analytic constituents are increasingly more dominant. So a graduated progression by itself is compatible with a significant analytic-synthetic thesis.⁷

⁵ The purely synthetic elements of Carnap’s construction are what we called above its “purely experiential elements”, i.e., elementary-experiences and the relation of recollection-of-similarity. Its purely analytic elements are conventional logico-mathematical and linguistic operations, like the operation of forming equivalence classes.

⁶ See Carnap (1928: 247–8). A more detailed account of the constructional stages of physical objects (and their concepts) appears in Carnap (175–211).

⁷ If the reader wonders what, then, is left of the disagreement between Carnap and Quine concerning the analytic-synthetic dichotomy, then three possible answers are: (a) According to Carnap, our body of knowledge is ultimately built out of purely experiential and purely conventional elements; according to Quine, no constituent of knowledge is either purely experiential or purely conventional. (b) According to Carnap, each element is uniquely divisible into experiential and conventional constituents; according to

Objection 2. The center-periphery methodology differs from the analytic-synthetic methodology with respect to what is arguably the most important principle of Quine's theory: the principle of universal revisability. The analytic-synthetic methodology views analytic statements as immune to revision, synthetic (empirical) statements as linked to a range of experiences whose occurrence would (under appropriate conditions) compel their revision. In contrast, the center-periphery methodology (like the NAS methodology) views no statement as immune to revision and the revision of no statement as compelled by any (specific) experiences. Given this major gap between the two methodologies, any similarity between them is of minor significance.

Response. The claim that the analytic-synthetic distinction is tied to the traditional revision principle is simply incorrect.⁸ This claim is refuted if not by the classical Kantian model, then by Carnap's model. In Carnap's model, analytic statements are based on linguistic conventions, and conventions in general are all too easy to revise.⁹ Furthermore, in Carnap's model even the revision of synthetic statements allows a significant degree of freedom. Given an infirming experience (or a series of infirming experiences) of a particular synthetic statement, Carnap allows a wide array of reactions. Thus, discussing a case in which a protocol sentence deduced from a physical theory is disconfirmed by experiment, Carnap says:

There are no established rules for the kind of change which must be made. For instance, the P[hysical]-rules can be altered in such a way that those particular primitive sentences are no longer valid; or the protocol-sentence can be taken as being non-valid; or again the L[ogical-linguistic]-rules which have been used in the deduction can also be changed (Carnap 1934: 317).¹⁰

Elsewhere, Carnap describes a hypothetical situation in which an empirical discovery might naturally lead to a conventional revision:¹¹

Quine, no element is uniquely divisible in this way. (c) According to Carnap, the application and satisfaction conditions of some elements are purely conventional while those of others purely factual (experiential); according to Quine, the application and satisfaction conditions of all elements are both partly factual and partly conventional.

⁸ Although Quine himself appears to endorse this claim (Quine 1951a: 43), it is wrong. (See also next fn.)

⁹ One might try to defend Quine's claim that for Carnap analytic statements are closed to revision by saying that Quine relates to *one* part of Carnap's dualistic model, namely *Science*. Quine would concede that linguistic conventions in Carnap's model are open to revision from the external perspective of Meta-Science, but his point is that they are not open to revision from the internal perspective of Science. (Since linguistic conventions are part of the *framework* of Science, they cannot be revised from within Science.) But to compare Quine's and Carnap's models we have to take into account their entire models, and in Carnap's case, his model includes both Science and Meta-Science.

¹⁰ The two sentences are transposed. In the original we have:

For instance, the P-rules can be altered in such a way that those particular primitive sentences are no longer valid; or the protocol-sentence can be taken as being non-valid; or again the L-rules which have been used in the deduction can also be changed. There are no established rules for the kind of change which must be made.

¹¹ For use of **boldface** within citations, see fn. 3 in Chapter 1.

If certain events allegedly observed in spiritualistic séances, e.g., a ball moving out of a sealed box, were confirmed beyond any reasonable doubt, it might seem advisable to [introduce a new conventional device, namely] **four** spatial coordinates (Carnap 1950: 212–3).

Furthermore, a conflict with experience, Carnap says, is even compatible with no revision at all:

[W]hen [a scientific hypothesis] proves to be L-incompatible with certain protocol-sentences, there always exists the possibility of maintaining the hypothesis and renouncing acknowledgement of the protocol-sentences (Carnap 1934: 318).

And lastly, Quine's principle of universal revisability is explicitly endorsed by Carnap:

[In "Two Dogmas"] Quine shows... that a scientist, who discovers a conflict between his observations and his theory and who is therefore compelled to make a readjustment somewhere in the total system of science, has much latitude with respect to the place where a change is to be made. In this procedure, no statement is immune to revision, not even the statements of logic and of mathematics.... With... this I am entirely in agreement (Carnap 1963: 921).

So, endorsing a principle of revisability that is both flexible and universal does not rule out endorsing a significant analytic-synthetic principle as well.¹²

Objection 3. There is no real conflict between CP and NAS since the two theses operate on different levels. While CP is a *descriptive, behaviorist* thesis, NAS is a *prescriptive, normative* thesis. CP describes scientists' habitual practices, including their tendency to regard certain statements as conventional, others as factual, and still others as occupying intermediate positions between the two extremes; NAS says that scientists' habitual practices are not inviolable principles of rational inquiry. Clearly, there is no conflict between the claim that scientists *normally act* in accordance with the analytic-synthetic distinction and the claim that scientists are not *normatively bound* by this behavior.

Response. This is an important objection whose validity depends on the role assigned to CP in Quine's argument. To identify this role, I propose the following analysis of the overall structure of "Two Dogmas of Empiricism". The paper is divided into two parts, sections 1 to 5, and section 6, respectively. The first part consists of a series of arguments to the effect that (i) the analytic-synthetic distinction is not well-defined, and (ii) the analytic-synthetic distinction is connected with a failed epistemology (namely, reductionism).¹³ None of these arguments demonstrates conclusively the unacceptability of the distinction. The first argument (cluster of arguments) merely shows that past attempts to define the distinction have failed,

¹² Once again, the agreement between Carnap and Quine on some aspects of revision leaves room for significant disagreements between them on other aspects. For example, Carnap might hold a strict bifurcation of types of revision (e.g., factual vs. conventional revision) to which Quine objects.

¹³ For exposition and critical discussion of these arguments see Grice and Strawson (1956), Gibson (1982), Gochet (1986), Boghossian (1996), Hylton (2007), and others.

not that any attempt to define it is bound to fail; the second argument shows that a popular epistemology incorporating the distinction suffers serious difficulties, not that any epistemology incorporating it is bound to suffer these difficulties or that these difficulties are in principle insurmountable. In spite of the inconclusive nature of these arguments, Quine arrives at a radical and uncompromising conclusion: The analytic-synthetic distinction should be altogether and in principle banished from philosophy.

Quine's "jump" from the claim that past attempts to establish the distinction face serious difficulties to the claim that the distinction is (in principle) an untenable dogma is obviously not a step in deductive reasoning. Rather, it is a methodological change of direction, similar to other methodological shifts in the history of philosophy, e.g., Kant's so-called "Copernican" revolution. We could, indeed, use roughly the same reasoning as Kant's (or the same skeleton of reasoning) to close the gap in Quine's argument:

Hitherto it has been assumed that all our knowledge must be reducible to ultimate elements of two kinds—analytic and synthetic. But all our attempts to account for knowledge based on this division have ended in failure. We must therefore make trial whether we may not have more success in the task of epistemology, if we suppose that knowledge is not divisible in this way.¹⁴

Quine's rejection of the analytic-synthetic dichotomy, like Kant's rejection of the traditional world-observer dichotomy, is a promissory note, to be made good by a new model of knowledge which demonstrates the virtues of the new methodology. The construction of such a model is undertaken in the second part (section 6) of "Two Dogmas". There Quine draws an outline of a new epistemic model, CP, whose task is to vindicate NAS. Quine is quite clear about the intended role of CP. Immediately following the introduction of the model, he presents it as implying the normative thesis NAS:

If this view [the view of knowledge represented by CP] is right, ... it becomes folly to seek a boundary between synthetic statements ... and analytic statements (Quine 1951a: 43).

Furthermore, tenets central to NAS follow from the model:

[I]t is misleading to speak of the empirical content of an individual statement ... Any statement can be held true come what may, if we make drastic enough adjustments elsewhere in the system. ... Conversely, by the same token, no statement is immune to revision (Quine: 43).¹⁵

¹⁴ Kant's original version is: "Hitherto it has been assumed that all our knowledge must conform to objects. But all attempts to extend our knowledge of objects by establishing something in regard to them *a priori*, by means of concepts, have, on this assumption, ended in failure. We must therefore make trial whether we may not have more success in the tasks of metaphysics, if we suppose that objects must conform to our knowledge" (Kant 1781/7: Bxvi).

¹⁵ The revisability principle, as we have seen above, is not unique to NAS, but it still constitutes an important component of NAS.

It is evident that Quine regards the center-periphery model as a *realization* of NAS. The model is designed to justify NAS by example, not by counter-example. CP and NAS are not two independent theses, the one descriptive, the other normative, the one capturing one aspect of our system of knowledge, the other, another. Rather, the two theses are designed to support one another: CP by providing an attractive picture of knowledge in accordance with NAS, NAS by motivating and articulating the normative principles underlying CP. Since NAS is a normative thesis, CP is assigned the role of a normative model exemplifying its principles.¹⁶

If this analysis is correct, the third objection is dissolved: CP is a normative thesis and its conflict with NAS is a genuine conflict, threatening the integrity of Quine's project. In the next section I will consider, and reject, two eliminative solutions to this conflict. In the last section I will present a third, constructive solution, which will bring us to the desired model of epistemic friction and freedom.

3.4 Dummett's Solution to the Inner Problem

One possible response to a conflict between two theses is to withdraw one of them. Thus, two available solutions to the present conflict are: (1) Eliminate NAS, and (2) Eliminate CP.

I. Elimination of the "no analytic-synthetic bifurcation" thesis (NAS)

This solution is recommended by Dummett. Quine's main achievement in "Two Dogmas of Empiricism", according to Dummett, is a new verificationist model of language and knowledge—CP—which, by being moderately holistic ("organic", in Dummett's terminology), avoids the pitfalls of earlier verificationist models. Since NAS threatens to undermine this achievement, Dummett recommends the elimination of NAS. Dummett finds support for his proposal in Quine's own change of heart in *Word and Object* (1960) and later writings, where, according to Dummett, Quine has "quietly dropped" NAS (Dummett 1973: 377).

Disputation. I will offer three grounds—exegetical, historical, and substantive—for rejecting this solution. The first, exegetical, ground is intended to show that this solution is incompatible with Quine's own comparative evaluation of, and degree of

¹⁶ One may try to argue for a descriptive reading of CP based on the third paragraph of section 6: "For vividness I have been speaking..." (Quine: 43–4). But even if it is natural to read a few sentences in this paragraph as suggesting a descriptive model, my argument is not affected. What I have attempted to do is *not* to show that no sentence in "Two Dogmas" is compatible with a descriptive model. Rather, I have attempted to show that (i) a descriptive model of knowledge would be extraneous to Quine's main argument in "Two Dogmas", (ii) the main argument of "Two Dogmas" calls for a normative model, and (iii) there is strong textual evidence to support a normative construal of the model. In contrast, interpreting CP as a descriptive model would raise difficult questions with respect to the interpretation of Quine's overall argument, the contribution of CP to this argument, and the role of the indicated passage in an otherwise normative section, sandwiched between two clearly normative passages. My tendency is to view the third paragraph as a digression, possibly foreshadowing certain trends of Quine's later epistemology.

commitment to, CP and NAS. The second, historical, ground, to which I will merely allude, is intended to show that from a broad historical perspective NAS is more important than CP. The third, substantive, ground supports NAS based on its own merits, and it partly explains why NAS should be incorporated in a model of knowledge of the kind we are seeking.

A. *Exegetical Reasons—Reasons Internal to Quine's Corpus—for Not Eliminating NAS:*

(a) *Early writings.* It is clear that prior to *Word and Object* (1960) Quine's espousal of NAS was full-hearted and uncompromising. Note the sharp tone and disparaging expressions in the following citations from "Two Dogmas":

[F]or all its a priori reasonableness, a boundary between analytic and synthetic statements simply has not been drawn. That there is such a distinction to be drawn at all is an unempirical **dogma** of empiricists, a metaphysical **article of faith** (Quine 1951a: 37).

[I]t is **nonsense**, and the **root of much nonsense**, to speak of a linguistic component and a factual component in the truth of any individual statement (Quine: 42).

If this view is right, . . . it becomes **folly** to seek a boundary between synthetic statements . . . and analytic statements (Quine: 43. Cited above).

In other articles from the same period, Quine's message is just as strong:

My misgivings over the notion of analyticity are . . . **misgivings in principle** (Quine 1953b: 139).

[A]nalyticity . . . is a **pseudo-concept** which philosophy **would be better off without** (Quine 1953c: 171).

In contrast, Quine's advocacy of the center-periphery duality is qualified: the idea of a distance from a center and a periphery is *merely* a "**metaphor**", brought in for the sake of "**vividness**" and *not to be taken quite literally* (Quine 1951a: 43). There is no doubt that faced with a choice between the two theses, Quine of "Two Dogmas" would have chosen NAS.¹⁷

(b) *Later writings.* In spite of endorsing a watered-down version of the analytic-synthetic distinction in some of his later writings, Quine never wavered in his rejection of the full-blooded distinction. Characterizing his own version of the distinction in *Word and Object*—stimulus analytic vs. stimulus synthetic—as a "behavioristic **ersatz**" (Quine 1960: 66) and a "**vegetarian imitation**" (Quine: 67) of the full-scale analytic-synthetic distinction, Quine declares that the full-scale distinction "only encourages **confused impressions** of how language relates to the world"

¹⁷ My understanding of what is central to Quine in "Two Dogmas" is different from Quine's own understanding in his "Retrospect" to this paper, written forty years later. There Quine approaches "Two Dogmas" in light of criticisms made during the forty years following its publication, and his perspective as well as concerns are different from those expressed in the original paper. (For Quine's focus in the 1991 retrospect see Chapter 6, Section 1, the end of Subsection I).

(Quine: 67). Then, commenting on the customary grouping of “synthetic truth”, “fact”, “objectivity”, and “accessibility to observation” together, he says: “So here that same analytic-synthetic dichotomy **intrudes** which we have found **so dubious**” (Quine: 247). And as usual, Quine is highly critical of Carnap’s uses of the distinction:

Carnap has long held that the questions of philosophy, when real at all, are questions of language... But why should this be true of the philosophical questions and not of theoretical questions generally? Such a distinction of status is of a piece with the notion of analyticity... and as **little to be trusted** (Quine: 271).

Quine’s impatience with “analyticity” is once again revealed in *Philosophy of Logic* (1970/86):

[I]t is time to **rein in our verbalism**. What are we trying to get at when we call a sentence analytic, or true purely by virtue of the language? (1970/86: 96).

In *The Roots of Reference* (1973) Quine continues to emphasize the difference between his own version of the distinction and Carnap’s:

[W]e have here **no such radical cleavage** between analytic and synthetic sentences as was called for by Carnap and other epistemologists (Quine: 80).

And as late as *Pursuit of Truth* (1990) Quine says:

The importance of analyticity for epistemology **lapses**, be it noted, in the light of holism (Quine: 55–6).

There is no question that Quine’s commitment to NAS persisted through the years.

B. *A Comparative Historical Consideration*. Historically, the analytic-synthetic distinction is tied to a long chain of philosophical dichotomies: the apriori vs. the aposteriori, reason vs. experience, reality vs. appearance, objectivity vs. subjectivity, realism vs. idealism, and so on. It is hard to exaggerate the role of these dichotomies in initiating philosophical movements. Rationalism, skepticism, empiricism, pragmatism, and positivism have all been driven by some or all of these dichotomies. In repudiating the analytic-synthetic distinction, NAS calls into question this entire line of philosophical bifurcations. From a historical perspective, therefore, NAS is a revolutionary thesis, challenging some of the most fundamental assumptions of philosophy and offering an altogether new approach to basic philosophical problems. In contrast, CP treads on well-known grounds, being a variant of the logical positivist account of knowledge and an offspring of classical empiricism. From the point of view of their critical as well as innovative potential, NAS is unquestionably of far greater philosophical interest than CP.

C. *Philosophical Strength of NAS*. This, in my view, is the main ground for not eliminating NAS. There are good philosophical reasons for preserving this thesis. These reasons are very different from the majority of reasons Quine himself gave in

support of this thesis in "Two Dogmas of Empiricism" (which have to do with language, meaning, mental entities, etc.), but they are still based on Quinean themes:

(a) *A Sound Methodological Principle.* A substantial epistemic consequence of the analytic-synthetic dichotomy is a rigid methodology of confirmation and infirmation: Synthetic truth is a matter of fact and is therefore confirmed or infirmed by whatever counts as factual evidence with respect to it; analytic truth is a matter of something other than fact, hence there is no such thing as factual evidence for (or against) it. This has direct ramifications for knowledge: units of knowledge are also units of truth, hence units of knowledge are also divided into synthetic and analytic. Synthetic units of knowledge are justified by appeal to the world; analytic units by appeal to X, where X is language, convention, pragmatic principles, concepts, mind, etc. *Synthetic units of knowledge are vulnerable to factual evidence, analytic units are immune.* This methodological principle, however, has never received a solid justification. Indeed, two important lessons that history has taught us contradict it. The first is that no unit of knowledge is immune to factual evidence; the second is that many revisions blur the line between the factual and the pragmatic. The second lesson was expressed by Quine in such passages as:

Revision even of the logical law of the excluded middle has been proposed as a means of **simplifying** quantum mechanics; and **what difference is there** in principle between such a shift and the shift whereby Kepler superseded Ptolemy, or Einstein Newton, or Darwin Aristotle? (Quine 1951a: 43).

The first lesson was elaborated by Putnam (1962, 1968). Putnam analyzed three episodes in the history of science: the replacement of Newton's definition of kinetic energy by Einstein's, the passage from Euclidean to non-Euclidean physical geometry, and the development of a non-standard logic for quantum mechanics. Whether one agrees with Putnam's specific analyses of these episodes or not, many of his points are methodologically sound. The revision of analytic units of knowledge (e.g., definitions, logical laws, and mathematical theories), Putnam notes, is not always conducted on a separate level of inquiry (e.g., the level of a conventionalist meta-theory) or guided by a special set of norms (e.g., purely pragmatic norms). Rather, in the course of scientific change, a definition may come to be treated as a law of nature (an identity statement),¹⁸ an abstract geometry as a physical theory (a theory of the

¹⁸ Putnam uses Einstein's definition of kinetic energy to exemplify the similarity between definitions and laws. The definition of kinetic energy, according to Putnam, is motivated by the same kind of considerations and has the same kind of role in Einstein's theory as such bona fide laws as "Moving clocks slow down" and "One cannot exceed the velocity of light" (Putnam 1962: 44–6).

The treatment of definitions as identity statements is also familiar from logical semantics: If T and T' are theories formulated within the framework of standard logic, and T differs from T' only in having " $a =_{\text{Df}} f(b)$ " where T' has " $a = f(b)$ " (or in having " $Ax =_{\text{Df}} Bx$ " where T' has " $(\forall x)(Ax \equiv Bx)$ "), then T and T' have exactly the same models, exactly the same logical consequences, and so on; i.e., T and T' are model-theoretically indistinguishable.

geometrical features of physical space), and a logic as a method of reasoning for various formal structures of objects (structures of objects possessing certain formal characteristics). In this way the laws of logic, for example, are in principle vulnerable to the world: if the formal behavior of objects conflicts (in some “deep” way) with a particular logic, then this logic is challenged by the behavior of those objects.¹⁹

Using the metaphor of a battle between Nature and Humanity to represent the pursuit of knowledge, we can describe the methodological rationale for NAS as follows: the analytic-synthetic dichotomy creates a false line of defense against Nature. Nature, the analytic-synthetic dichotomy induces us to believe, is in principle incapable of threatening the analytic zone of our knowledge. But Nature (or reality) might—and some say, did—find ways to encroach upon this zone. The analytic-synthetic policy of complacency in the analytic zone, careful measures for establishing the correctness of our theories in the synthetic zone, is therefore unwarranted. We do not know in advance where Nature will choose to strike next, and by restricting our defenses to the synthetic front of knowledge we might be creating a Maginot line of defense. As a matter of prudent strategy, it is therefore incumbent upon us to maximize the maneuverability of our cognitive resources, and just this is accomplished by NAS.

Now, one possible objection to this line of reasoning can be formulated as follows: granting the unpredictability of Nature, it is still not clear how Nature can surprise us when it comes to our own concepts. If the meaning of “A” includes “B”, what can Nature do about it? How can Nature falsify our judgment that every A is a B? This challenge can be expressed both in Kantian and in positivist terms: in Kantian terms the claim is that some truths are grounded exclusively in meaning; the positivist claim is that some statements are true by postulation. This leads to our second point.

(b) *The Genetic Fallacy Argument.* Quine aptly responded to the positivist challenge in “Carnap and Logical Truth” (1954a). There Quine argues that to view a statement as true by postulation is to commit a genetic fallacy.

The genetic fallacy is, traditionally, a fallacy of relevance: the genesis of a statement Z has nothing to do with the rationality of upholding Z. Quine’s genetic-fallacy argument proceeds as follows: suppose a statement Z made it into our corpus of knowledge on grounds of analyticity. What is the significance of this fact? This fact shows that at a particular moment in the history of our corpus Z was accepted based on postulation. But this fact has nothing to do with the reasons for maintaining (or readmitting) Z at a later stage in the development of our corpus. Analyticity and

¹⁹ This way of looking at logic is common among quantum logicians. The rules of classical logic, these logicians argue, are “obeyed” by systems of objects embeddable in a Boolean structure, but systems exhibiting other formal structures, for example, the structure of a non-distributive complemented lattice, require a different logic. See Putnam (1968), Finkelstein (1969), Hughes (1989, Chapter 7), and others. This way of looking at logic may also underlie Tarski’s statement that “certain...experiences of a very fundamental nature may make us inclined to change just some axioms of logic” (Tarski 1944a: 31–2). (Tarski’s view is further discussed in Chapter 10, Section 2.)

syntheticity are no more than genealogical properties of statements—significant *historically*, but not *normatively*. Normatively, our task is to develop a *correct* theory of the world. Eventually, our postulates must earn their place in our theories not just in virtue of being convenient but also in virtue of being true.²⁰

A similar argument was presented by J. G. Maaß (1789) against Kant's version of the analyticity claim. Whether a given statement satisfies Kant's conditions on analyticity varies, according to Maaß, from speaker to speaker and from one occasion of speech to another: some speakers arrive at "A is B" by being taught that the concept B is contained in the concept A, others, by observation; some speakers *extend* their knowledge by adding "A is B" to their corpus, others (or the same speakers on other occasions) do not. Kant (through a review by Schulze)²¹ rejected Maaß's criticism on the ground that relations between concepts are *fixed*: we can change the reference of our words, but not the reference of our concepts. Quine, however, denies the "fixity" of concepts. The non-fixity of concepts is our third reason for supporting NAS.

(c) *The Non-Fixity of Concepts*. The traditional analytic-synthetic distinction involves the assumption that our concepts are fixed: there is a fixed fact of the matter concerning the content (meaning, reference, composition) of our concepts; hence there are immutable relations between them. Theories change, concepts are permanent. Against this view Quine introduces his principle of the non-fixity of concepts. This principle is closely related to his thesis of the inseparability of language from theory, according to which the meaning of terms is inextricably tied with theories in which they occur: the meaning of "atom" with theories of matter, the meaning of "light" with optical theories, etc.²² It follows from this thesis that change in theory involves change in the content of our concepts; that is, the content of our concepts is not fixed.

The thesis of the non-fixity of concepts is also elaborated by Putnam (1962). Putnam argues for the non-fixity of concepts on two grounds: (i) a Wittgensteinian analysis of certain scientific concepts as (in effect) "family resemblance" concepts, and (ii) the observation that sometimes scientific concepts emerge from a radical

²⁰ The requirement of establishing the truth of postulates is commonplace, according to Quine: any theoretical hypothesis is, at the time it is made, no more than a conventional postulate (being justified by nothing more than "the elegance or convenience... [it] brings to the containing body of laws and data"), and therefore the process whereby we test it for agreement with experience is just the process whereby a conventional postulate is tested for truth (Quine 1954a: 121).

The need for a veridical justification of statements that were originally introduced as postulates arises also in mathematics, where postulates are often thought of as *obviously* true. Thus, Cantor, for example, started by holding the well-ordering principle a self-evident postulate, but later required a rigorous proof of this principle (see Moore 1982, sections 1.5 and 1.6).

²¹ Review of the second volume of the *Philosophisches Magazin* (mentioned in Allison 1973: 174–5).

²² Even in literature our use of terms is affected by theory, according to Umberto Eco. Eco brings the following example: "The bible obliges its readers to interpret a whale as a big fish; contemporary novels presuppose that whales are to be interpreted as mammals. When Melville seeks to play a *double jeu* and to tell of a whale that must be considered at the same time as both a fish and a mammal, he spends an entire chapter making his readers eager to overlap two different encyclopedias" (Eco 1995: 24).

change in theory changed in content but unchanged in identity. Explanation: Putnam observes that some scientific concepts are interwoven in our theories in such a way that their content is determined by the cluster of laws associated with those theories. (The concept of energy, for example, is embedded in a cluster of laws containing the law of conservation of energy, laws concerning chemical, gravitational, and nuclear energy, laws relating energy to mass and momentum, etc.) A change in some of these laws results in a change in the content of those concepts, but since those concepts are connected with a host of other laws (which presumably do not all change at once), their identity is preserved across time. Whether and in what way a given law-cluster concept will change its content is an open question, according to Putnam: nothing in the concept itself can protect it against change. Putnam restricts his analysis to certain “privileged” concepts of physical science: “energy”, “light”, “straight line”, etc. But (as Putnam is well aware) there is nothing in his analysis to exclude other concepts from exhibiting the same behavior. Indeed, it follows from the genetic fallacy argument that even concepts originally introduced by convention can become interwoven in our theories in the way Putnam describes. (Putnam claims that at every stage in the development of science some concepts must play the role of “fixed points”. This claim granted, it does not follow that the *same concepts* must play this role at every stage, i.e., that there are *inherently* “fixed” concepts.) The view that all concepts are, in principle, sensitive to change in theory is tantamount to Quine’s thesis of the inseparability of language from theory. It follows from this claim that no truth is guaranteed to be true by the “fixed” meaning of its terms; that is, no unit of knowledge is analytic in the traditional sense.²³

These, then, are my reasons for rejecting Dummett’s proposal of eliminating NAS: (A) Exegetically, Quine is strongly committed to NAS; (B) historically, NAS is a revolutionary thesis of much philosophical interest; and most importantly, (C) substantively, NAS is supported by three sound principles: (a) the principle of maximizing the maneuverability of our cognitive resources, (b) the principle of the normative insignificance of the genesis of our theories, and (c) the principle of the non-fixity of concepts.²⁴

²³ I should note that although I disagree with Putnam’s view that some concepts are *inherently* “fixed”, my overall view of what is important in Quine’s theory is, as my earlier references clearly indicate, informed by much of what Putnam says in his (1962) paper. This includes those parts where he sets out to show that “Quine is overwhelmingly right in his critique of what other philosophers have *done* with the analytic-synthetic distinction”, though not those parts where he tries to show that “Quine is wrong in his literal thesis” about the analytic-synthetic distinction (Putnam: 42). It should also be noted that Putnam himself attributes greater importance to the way Quine is right than to the way he is wrong. Thus, while Putnam adamantly says “That Quine is wrong I have no doubt”, he insists that “in a deeper sense . . . Quine is right” (Putnam: 36).

²⁴ It is important to emphasize that this view is not tantamount to meaning skepticism. In the same way that viewing theories as changeable (revisable) is not tantamount to theory-skepticism, so the changeability of concepts does not amount to meaning-skepticism. Take physics, for example. During each stage in the development of physics, physical theories have a distinct identity; similarly, during each stage in the development of physics, physical concepts have a distinct identity.

II. Elimination of the “center-periphery” thesis (CP)

This is the second eliminative solution to the tension between CP and NAS: If NAS is worth preserving and CP is incompatible with NAS, why don't we eliminate CP?

Disputation. We cannot eliminate CP *tout court*, since some positive model of knowledge is needed to validate NAS. Moreover, CP succeeds in making two indispensable contributions to Quine's model: (A) It gives structure to its holistic principles. (B) It provides it with a means for linking knowledge to reality.

A. *Structuring Holism.* Without something like the center-periphery duality, Quine's system of knowledge is a formless body of haphazardly interrelated statements. Such an amorphous holism is, as we have seen in Chapter 2, Section 2, evidently unacceptable. First, it lacks explanatory power: to give an explanatory account of knowledge is to impose a meaningful structure on its elements. Second, it renders our body of knowledge utterly unmanageable: a structureless object can be grasped either in its entirety, hence in one act of cognition, or not at all; but our system of knowledge is clearly too large, too complex, and too open-ended to be grasped in one act. Conclusion: A viable holism must impose a manageable structure on our body of knowledge; i.e., a viable holism is a structured holism and one whose structure is in principle manageable.

CP introduces both structure and simplicity into Quine's holistic model. It postulates a relatively simple structure of knowledge: a field or a sphere with two distinguished zones. The one zone—*center*—is the center of interconnectedness, the other zone—*periphery*—marks the outer reaches of the system. From the center our body of knowledge stretches outward towards the periphery, moving gradually from the universal to the particular, from the general to the specific, from the abstract to the concrete, from the conventional to the factual, until reaching its outer boundaries. The periphery is linked to the center by a network of cognitive routes (routes of deductive and statistical inference, routes of abstraction and generalization, routes of application and instantiation, routes of revision, explanation, justification, etc.), and it is through this network of routes that information, evidence, and other epistemic commodities move across the system.

B. *Anchoring Knowledge in Reality.* By itself, Quine's holistic conception of knowledge, with its principles of interconnectedness, universal revisability, NAS, pragmatism, etc., is compatible with a coherentist or even an idealist epistemology. Any system of knowledge will satisfy these principles, provided certain conditions having to do with its internal structure are satisfied. The transformation of Quine's holistic model into a model of knowledge of reality is accomplished by CP. CP delineates the inner

It is also important to indicate that the above considerations in support of NAS are very different from those criticized by such authors as Boghossian (1996), as will be noted in Chapter 6, Section 2. Nevertheless, one aspect of the analytic-synthetic distinction which is right, in my view, is incorporated into the present model. See discussion of Grice and Strawson in Chapter 6, Section 2.

structure of our system of knowledge in relation to reality. It anchors our corpus of knowledge in reality on the one hand, and in the mind on the other. Our corpus is linked to reality through the periphery, and the impact of reality permeates the corpus due to interconnections between its elements. These interconnections are regulated in the center, where the influence of language, method, conceptual frameworks, pragmatic norms, etc. is most strongly felt. Changes may take place anywhere in the sphere, but all changes must square with reality (directly or indirectly) through the periphery. Even highly abstract physical theories must square with experience through the periphery, e.g., through direct (or indirect) experiments. If we call the requirement that our system of knowledge be anchored in reality *realism* (or *basic realism*)²⁵, then a model of knowledge based on NAS requires something like CP to render it (basically) realistic.

In making these important contributions to Quine's non-traditional model, however, CP turns into a super-traditional model: one fixed center—the site of logic, mathematics, and philosophy—surrounded by scientific theories whose supporting evidence (in the form of experiential reports) is transported from the periphery. We may liken this structure to a traditional industrial economy: one metropolis—the locus of capital, services, means of communication—surrounded by industrial sites whose raw materials are transported from quarries and mines across the country. The system as a whole is interconnected, but *the norms governing its center and periphery are radically different*. The norms governing the center (“service sector”) are *pragmatic*, those governing the periphery (“productive sector”) are *factual* or *veridical*. And *it is this rigid division of the cognitive norms that is irreconcilable with NAS*.

3.5 A New Solution

Our analysis suggests that NAS is an innovative thesis of much philosophical interest, while CP makes an important contribution to demonstrating its epistemic viability. Neither thesis should be eliminated, yet the two stand in a fundamental tension. One natural way to solve the problem is to reconfigure one (or both) of these theses by identifying those elements which are responsible for the conflict and replacing them by others that will institute agreement. In view of the greater importance (promise) of NAS, it seems reasonable to approach the problem as one of reconfiguring CP.

My solution is based on the following observation. There is no conflict between the idea, systematized by CP, that the two main sources of knowledge are mind and world, and the idea, central to NAS, that the division of labor between them within any unit of knowledge is neither fixed nor lasting. Rather, the root of the conflict is the assumption, implicit in CP, that each statement, and, by extension, each theory, is

²⁵ *Basic realism* will be explained and discussed in Chapter 5, Section 4.

located within a fixed region (if not a fixed point) of the sphere—within a fixed distance (or limited range of distances) from its center and boundary—and this distance is determined by the fixed ratio of factual to pragmatic (conceptual) factors in its truth and justification. The presumption of a fixed location (or a limited range of locations) involves the requirement that a fixed range of factual (for Quine, experiential) input and a fixed range of conceptual (for Quine, pragmatic) input be associated with each statement (theory), but just this is contested by NAS. It is thus not the existence of two distinct sources of knowledge, metaphorically referred to as “center” and “periphery”, that is the root of the conflict; rather, it is the postulation of fixed factual conditions and fixed conceptual (conventional, pragmatic) conditions for the truth and justification of each statement that is responsible for it. The fixity requirement has two dimensions: (a) At any given point in the duration of our corpus each statement (theory) occupies a fixed (unique, absolute) position within the sphere; and (b) while in the course of history a statement (theory) may change its exact location (e.g., as a result of a re-axiomatization of a theory to which it belongs), its general area as well as its rough distance from the center and the periphery remain the same. This is true especially of observation statements, which always occupy the periphery, and of logical statements, which are always in the center.

My solution to the inner conflict in Quine's theory lies in replacing this *traditional*, *static*, and *absolutist* model by a new, *dynamic* model. There are still two zones, center and periphery, corresponding to the two basic constituents of human knowledge, mind and world, but no fixed distribution of statements and theories within the model. As a theory of the formal laws governing structures of objects—for example, the laws relating identity to cardinality—logic is in the periphery; as a theory of the logical laws governing statements and sets of statements, logic is in the center. As a theory saying that a structure whose objects are all identical is a structure of one object, logic is about the world; as a theory saying that the statement “all objects are identical” implies the statement “there is exactly one object”, logic is about language. As a theory of the formal structure of “real” objects (configurations of objects), logic is anchored in reality; as a theory of the formal structure of our thought of objects, logic is embedded in the mind. A similar duality exists in other disciplines. Take zoology, for example. As a theory of the hierarchy of life, zoology is in the periphery; as a system of concepts of life (as a hierarchy of zoological concepts), zoology is in the center.

Normatively, the project of knowledge consists of two tasks: the task of constructing a serviceable set of conceptual tools, and the task of constructing a correct and informative theory of the world. Accordingly, our system of knowledge is measured by standards of two types: veridical standards—truth, confirmation, prediction, explanation; and pragmatic standards—economy, convenience, utility. Center and periphery are now distinguished not extensionally (since all theories are charged with satisfying both types of standard), but functionally, namely, by the different tasks or

standards attended to while we, or our theories, reside in each region.²⁶ In constructing a system of knowledge, we move *back and forth* between the center and the periphery: we set out to improve the overall usefulness of our conceptual apparatus, and we set out to verify the correctness of our theories. Each theory in our corpus must square off with some facet of reality: logic with the formal behavior of objects, physics with their physical behavior. Given a conflict with reality, we are given great latitude in devising a solution (*epistemic freedom*). We may resolve the conflict by making changes in our physical theory of the world, keeping logic “fixed” in the background; or we may resolve it by changing our theory of the formal structure of the world, keeping the physical laws “fixed” in the background. We may resolve the conflict by placing physics in the “front” (periphery), logic in the “rear” (center); or we may resolve it by placing logic in the front, physics in the center. By holding our notion of physical object fixed, we constrain the variability of our logic; by holding the laws of logic fixed, we restrict the variability of our physics.²⁷ Ontology plays a

²⁶ This is our response to one of Dummett’s criticisms of Quine’s model, which challenges the difference between center and periphery in that model:

In accordance with [NAS], the revision of truth-assignments to the sentences of the language which is elicited in response to a recalcitrant experience may not affect any of the peripheral sentences, but only those lying below the periphery. But, if this is so, then, it seems, experience does not impinge particularly at the periphery; rather, it impinges on the articulated structure of our language as a whole, not at any one particular point. **In that case, it becomes difficult to see how we can any longer maintain a distinction between periphery and interior:** the periphery was introduced as that part of the structure at which the impact of experience is first felt (Dummett 1973: 376–7).

According to the changes we institute in Quine’s model, the difference between center and periphery concerns the type of standards employed in a given revision. A recalcitrant experience may be handled either in the center or in the periphery (or in the intermediate zone). It is handled in the periphery when the revision that takes care of it is largely factual (or based largely on factual considerations) and in the center when the revision is largely pragmatic/conceptual (or based largely on pragmatic/conceptual considerations). But both types of revision might involve any discipline. A recalcitrant experience in say, physics, might lead us to make either a factual or a conceptual revision (or both), and either revision might involve either (or both) physics or (and) some other disciplines, e.g. logic.

²⁷ By speaking about this pattern I do not rule out other patterns, including patterns involving joint changes in logic and physics. In such cases, some elements of logic, physics, and everyday rationality are held fixed, while others are changing. The factuality of logic will be systematically discussed in Chapters 9 and 10. This discussion will also explain why another criticism by Dummett, one concerning the claim that logic is not analytic, does not affect our own solution to the inner tension in Quine’s model (whether this criticism applies to Quine’s original model or not). Dummett says:

For language to be an articulated structure, there must be links between the sentences: it is the presence of these links which constitute the totality of sentences as a structure, that determine the position of each sentence within the whole. We should naturally take these links to consist of inferential connections (perhaps including inductive as well as deductive ones). For inferential connections of a deductive character, at any rate, we know of no way of formulating the existence of such connections that will not have, at least as a by-product, the effect of guaranteeing the truth of certain sentences, which will be precisely the analytic ones. Hence a thesis which denies the existence of analytic sentences calls in question the existence of any inferential links between non-analytic sentences (Dummett: 377).

In Chapter 10 we will see that there is a way of “formulating the existence of [deductive] connections” that will not render them analytic. Indeed, Dummett himself is not sure that such a way does not exist. He says:

similar role, as do other branches of philosophy. Philosophy, on this picture, is not above (or below) the sciences; rather philosophy—through its theories of logic, object, truth, rationality, and the like—faces the court of knowledge both in the center (i.e., on the level of concepts and method), and in the periphery (i.e., on the level of fact), just like other sciences. And like other sciences, philosophy has its own identity, methods of inquiry, and relative autonomy: philosophy is not reducible to empirical psychology.

While Quine's narrow empiricism (e.g., his conception of the periphery as purely experiential) is rejected by this model (which allows logic to lie in the periphery without becoming an experiential discipline), many of his theses are integrated into it: NAS, the interconnectedness of knowledge, the universal revisability of beliefs, the inseparability of language from theory, basic realism, scientific empiricism, anti-foundationalism, and certain important features of the center-periphery duality. In particular, the new center-periphery structure continues to make the two main contributions of the original one: (i) structuring the interconnections between statements (theories), (ii) linking our system of knowledge to reality. Unlike the old model, however, the new model is thoroughly dynamic: at each stage in the development of our corpus, each science lies in the periphery in some contexts, in the center in others; and in the course of history, sometimes one constellation of theories, sometimes another, occupies a more prominent place in the periphery/center.

This is my solution to the inner conflict in Quine's epistemology.²⁸ Turning back to the principles of friction and freedom, the foundational-holistic methodology, and the "building blocks" delineated in the beginning of this chapter, we see that this solution leads to a model that exemplifies, at least in broad strokes, these fundamentals. In the next three chapters I will further develop this model and clarify its nature. In Chapter 4, I will expand my account of its dynamic structure, in Chapter 5, I will discuss some of its central principles, and in Chapter 6, I will further clarify its differences from Quine's model and show how it is immune to major criticisms of the latter.

One possible objection to the present model rests on a historico-empirical argument. According to this argument, the view that logic and mathematics must be able to occupy the periphery at some points in the historical development of our system of

"It would take a much longer enquiry than will be pursued here to **determine whether there is an actual inconsistency** [between affirmation of deductive connection and rejection of analyticity]" (Dummett: 377).

²⁸ I should note that around the same time that my dynamic model of knowledge was first delineated in print (Sher 1997a, 1999a) Michael Friedman presented a dynamic conception of knowledge (Friedman 2001, based on lectures given in 1999). For the same reasons as those noted in Chapter 2, Section 2, fn. 5, I will not be able to engage in a comparison of my dynamic conception of knowledge and Friedman's in this work, but in Chapter 5, Section 3, I will note the similarity in our rejection of the "pure" apriori and acceptance of a weaker version of the apriori (Friedman's "relativized apriori" and my "quasi-apriori"), and in Chapter 6, Section 2, I will explain why Friedman's criticism of Quine's holism does not apply to my foundational holism.

knowledge is undermined by the fact that there have never been real revisions of logic or arithmetic; there have only been minor corrections of earlier misconceptions. In response, let me note that historically, there were cases of revision in mathematics and logic, including rather dramatic cases, many of which took the form of, or were based on, *factual discoveries*. One of the most dramatic revolutions in our mathematical thinking, a revolution that shook mathematicians, physicists, philosophers, and laypeople alike, was the discovery of non-Euclidean geometries in the nineteenth century and the rejection of Euclidean geometry as the geometry of physical space by Einsteinian relativity. Other historical examples include revisions of our firm convictions that Aristotelian logic is the ultimate logic, that the ontology of mathematics is finite (or, later on, at most denumerable), that mathematics is complete (all mathematical truths are provable), etc. These changes led to significant revisions in the disciplines involved as well as in related disciplines (e.g., physics in the case of geometry).²⁹

I should emphasize, however, that my main argument for the view that logic and mathematics must be able to reside in the periphery is not historical. My argument is philosophical in nature, starting with the universal friction requirements and ending with reasons why, in order to perform its designated tasks in our system of knowledge, logic must be grounded not just in the mind but also in the world. The first step of this argument was given in Chapter 1, Section 2, the second in the present chapter. The rest will be given in a series of discussions, starting from clarifications and examples of logic's and mathematics' presence in the periphery (next chapter), and proceeding through discussions of reality and realism (Chapter 5, Sections 2 and 4), a new correspondence account of mathematical truth (Chapter 8, Section 4), and explanations of logic's grounding in reality (Chapter 10), which includes discussions of (i) why logic requires a grounding in the world, (ii) what specific features of the world logic is grounded in, (iii) how this grounding enables it to perform its role in knowledge, (iv) how it explains its relative immunity to refutation, and (v) why, nevertheless, it is important to recognize and understand the possibility of error and revision in logic.

Let me conclude with an issue that some readers might wonder about: the relation between the model I delineate in this essay and Quine's original model. Different readers might view this relation differently, depending on their understanding of Quine and their judgment about what is central to his theory. Some may view the present model as a bona fide Quinean model, others as an anti-Quinean model; some

²⁹ Of course, each of our examples of a revolutionary change is controversial (for example, some still believe that mathematical ontology is finite). But historically, these examples are significant. Note that the fact that there have not been revisions in individual logical beliefs such as the belief in modus ponens (or instances thereof) is just as insignificant as the fact that there have not been revisions in such individual physical beliefs as the belief that grass is green or that snow is cold (or even the more general belief that if an object is all red then it is not yellow).

may view it as a “neo-Quinean” model (as I did in Sher 2010), others as a model that combines Quinean and non-Quinean elements (as I do today). All this is unimportant for the present project. What *is* important is the model itself and its major principles.³⁰

³⁰ Thus, when in the rest of this essay I refer to the model by “the present model”, “my/our model”, “the new model”, or “the/our dynamic model”, these modes of reference are to be understood as tags rather than as definite descriptions. In Chapter 4, I will focus on points that, in some cases, enhance the similarity between the present model and Quine’s model; in Chapter 5, on issues that are not specifically connected to Quine; and in Chapter 6, on differences between some of Quine’s views and mine, criticisms of some of his views, and explanations of why the present model is largely immune to influential criticisms of his model. But even there, the goal is to prevent misunderstandings of, and demonstrate the viability of, the model I support, rather than to offer a scholarly account and evaluation of Quine’s writings.

4

Dynamic Model—Two Dimensions of Change

To clarify the nature, scope, and ramifications of the new model, I will offer a sketch of its dynamic structure. The sketch will center on the two dimensions along which concepts, statements, and theories shift their position within the model: the *time dimension* and the *context dimension*. The sketch will focus on (i) the movement of statements and theories along the two dimensions, (ii) the sense in which logic and philosophy lie, in certain respects and at certain times, in the periphery, (iii) the normative principles (associated with the *no analytic-synthetic bifurcation* thesis (NAS)) underlying the model, and (iv) some of the model's ramifications for outstanding philosophical problems. In drawing the sketch I am aiming at clarity and, to a certain extent, systematicity, but not at a comprehensive or exhaustive account. The sketch draws both on Quinean and on non-Quinean themes. The center-periphery metaphor is broadly interpreted, not just in the sense that the center and periphery represent the interface of our system of knowledge with mind and reality in general and not only with pragmatic conventions and sensory experience, but also in the sense that the center-periphery metaphor represents a *family* of dualities: “mind and world”, “concept and object”, “rear and front of the battle for knowledge”, and so forth.

Two Dimensions of Change. The most distinctive characteristic of the new model, compared with the old one, is the *movement* of cognitive elements within the model, in particular their movement from the center all the way to the periphery and vice versa. This movement typically takes place along two axes, *time* and *context*. Diachronically, the new model is dynamic in the sense that as our body of knowledge develops, the deployment of concepts, statements, and theories within it changes in response to changing circumstances. Synchronically, it is dynamic in the sense that at each stage in the development of our system of knowledge each statement, concept, and theory falls within a multiplicity of contexts, and represented from the varied perspectives of these contexts, its position within the model changes.

4.1 Contextual Dynamic

We can distinguish two sources of contextual dynamics: (1) *multiplicity of perspectives*, and (2) *multiplicity of factors*.

I. *Change of perspective*

Change in perspective may involve change in orientation. Viewed from one perspective, a concept, statement, or theory is world-oriented, viewed from another, language- or mind-oriented; viewed from one perspective, a given subject matter is factual, viewed from another, mental or conceptual. Two sources of such change in perspective are (A) *the two-sidedness of language*, and (B) *the plurality of human interests*.

A. *The Two Faces of Language (Semantic Ascent and Objectual Descent)*. One source of the multiplicity of perspectives is the two-sidedness of language. It is an old observation (made by medieval philosophers) that sentences in general can be read both in the linguistic mode and in the factual mode. Take, for example, the sentence “Whales are mammals”. Read in the linguistic mode, this sentence says that “whale” is subsumed under “mammal”; read in the factual mode, it says that whales possess the characteristic properties of mammals. Read in one mode this sentence is language-oriented, read in the other, world-oriented. One important vehicle of such change in orientation is the truth predicate. The truth predicate enables us to speak about the world by speaking about language, and to speak (in certain ways) about language by speaking about the world. Instead of saying that snow is white, we can say that “Snow is white” is true, and instead of saying that “Snow is white” is true, we can say that snow is white. Quine introduced the term “semantic ascent” for one direction of this shift; I will use “objectual descent” for the other direction.

(a) *Truth is both language- and world-oriented*. The dual principle of semantic ascent—objectual descent is encapsulated in the “truth schema”. But the truth schema itself, as noted by many (including Tarski and Quine), allows both a linguistic and a factual reading. Read in the linguistic mode, the truth schema is a disquotational schema; read in the factual mode, a correspondence schema. Under one reading, truth has to do with an intralinguistic relation; under the other, with a relation between language and the world. While the two readings yield the same T-sentences, their explanation of truth is radically different. The disquotational account of truth is linguistic, the correspondence account is metaphysical; the disquotational notion of truth is language-oriented, the metaphysical notion is both language- and world-oriented.

(b) *Objectual descent: logic is world-oriented*. The principle of objectual descent provides one explanation of logic’s presence in the periphery. It follows from this principle that the logical laws, despite their appearance (e.g., in the case of sentential logic) as laws governing the behavior of language, reflect certain regularities in the behavior of the world. Thus Quine comments:¹

¹ For **boldface** within citations see fn. 3 in Chapter 1.

Logical theory, despite its heavy dependence on talk of language, is... **world-oriented** rather than language-oriented (Quine 1970/86: 97).²

And elsewhere he says:

Consider... the **logical truth** 'Everything is self-identical', or ' $(x)(x=x)$ '. We *can* say that it depends for its truth on traits of the language (specifically on the usage of '='), and not on traits of its subject matter; but we can also say, alternatively, that it **depends on**... [a] **trait, viz., self-identity, of its subject matter, viz., everything** (Quine 1954a: 113).

On the second interpretation, the law of self-identity is world-oriented.³

(c) *The epistemic priority of the factual: the immanence of knowledge.* While statements and theories within the model are both fact- and language-oriented, epistemically their factual orientation takes priority over their linguistic orientation. Using a Quinean figure of speech, we can call this the thesis of the "immanence" of knowledge. This thesis says that we know the world through our theories (i.e., from a standpoint internal to our theories), but speaking from within a theory is speaking factually. We use words in theories to speak about objects, and in so doing we give priority to their factual reading. (We can also view this as the principle of "use over mention".) The epistemic priority of the factual over the linguistic means that all knowledge is, in a certain basic sense, fact- or periphery-oriented. In Quine's case, this view is reflected in his comment, mentioned earlier, that logic is "world-oriented rather than language-oriented", as well as in his comment that all branches of knowledge, including philosophy, are about "what there is" (Quine 1960: 275).⁴

(d) *Philosophy is fact-oriented.* The immanence of knowledge means that all sciences are fact-oriented. We can explain the way this principle applies to philosophy (even beyond logic and ontology) by contrasting the factuality of philosophy with its alleged conventionality. Consider Kant's epistemology once again. Kant's declared motivation for his "Copernican revolution" is largely pragmatic: all attempts to establish knowledge based on the traditional conception of the relationship between the knower and the world have failed; let us see whether by turning this relationship around we will not be able to arrive at a better theory.... But the theory constructed based on this move makes *factual* claims about its subject matter, viz., human knowledge. Speaking in terms of aims of theory, we can say that Kant's theory aims not just at an *expedient* account of human knowledge, but also (and indeed primarily) at a *correct* account. Kant's goal is to identify the *real* conditions for the

² Quine goes on to say that it is "the truth predicate" that "makes" logic world-oriented. This is a point he already made: "[T]he truth predicate is... doing an active job of separating logic from language" (Quine 1970/86: 97). We will get back to this in Chapter 10, Section 3.

³ These comments, however, are somewhat compromised in light of Quine's empiricism, as we will see in Chapter 6, Section 1.

⁴ I should note, though, that in light of Quine's empiricism, my understanding of the factuality of philosophy is somewhat different from his, as we will presently see. (See also discussion in Chapter 6, Section 1.)

possibility of knowledge (not some convenient pseudo-conditions), to give a *correct* analysis of the basic structure of our cognition (not a convenient fiction about its structure), to provide a *valid* proof of the possibility of human knowledge (not the semblance of a valid proof, i.e., a chain of statements that can conveniently be imagined to be a valid proof). Kant aims at a *correct* theory of the structure and possibility of human knowledge, and in this sense his theory, and philosophy more generally, is oriented toward reality. This is perhaps the deepest sense in which philosophy lies in the periphery. Philosophy seeks to obtain an understanding of knowledge, truth, morality, valid inference, and so on, as they in fact are: their *true* factors, *true* structure, *true* possibility, *true* difficulties—all as they *truly* or *factually* are, not as we may conveniently wish, or imagine, or postulate them to be.

Critical Note: The factuality of philosophy means not that philosophy is less pragmatic than other sciences, but rather that the norms of truth (accuracy, absence of error, etc.) are as central to philosophy as they are to other sciences.

B. *Multiplicity of Interests.* Another important source of the multiplicity of perspectives is the multiplicity of interests. The multiplicity of interests leads to the development of multiple theories of the same subject matter. Take the subject of correct reasoning, for example. By studying correct reasoning from a purely formal and theoretical standpoint, we arrive at a “pure” logic, while by studying it from a point of view encompassing also the material and the practical, we arrive at a (general) theory of rationality.⁵ Differences in interest lead to differences in perspective which, in turn, may lead to radical differences in our conception of a given subject matter. By thinking of logical laws as laws governing the interaction of physical properties (states), quantum logicians have come to view logical laws as factual. By thinking of these laws as governing the manipulation of symbols, the formalists came to view them as conventional. Differences in approach to a given subject matter may either agree or disagree with one another: the intuitionistic approach to logic competes with the classical approach, whereas the modal approach leads to its extension. Different theories of the same subject matter are located in different areas of the field: a veridical theory of logical validity lies closer to the periphery, a pragmatic theory of logical validity lies closer to the center.

Critical Note: Not every point of view on every subject matter is validated by the model. That is, the model does not represent a methodology of “anything goes”. On the contrary. It is a central principle of the new model that every theory is subject both to veridical norms (truth, accuracy, avoidance of error, justification, etc.) and to pragmatic and other extra-veridical norms (simplicity, unity, generality, fruitfulness, etc.). This means that in the new model the standards for the acceptance of theories are *higher* than in most other models. We can express this by saying that each theory

⁵ Harman (1986, 2009) offers yet another perspective on the relation between logic and the theory of rationality. See the subsection on the normativity of logic in Chapter 10, Section 6.

is subject not just to the norms of utility but also to the norms of truth, and not just to the norms of truth but also to the norms of utility. As a result, most points of view on a given subject matter are rejected by the model either as conducive to error or as inexpedient, unfruitful, etc. Nevertheless, the principle of a multiplicity of viewpoints is affirmed by the model.

II. *Multiplicity of factors*

The contextual mobility and malleability of subject matters, statements, and theories within the model is not, primarily, a subjective matter, having to do with our personal choices, whims, and proclivities, but (to a large extent) an objective principle, having to do with the multifariousness of the world and the plurality of factors constraining the knower.

A. *Multiplicity of Factors Constraining the Knower (Knowledge is Not a Mirror Representation of Reality)*⁶. That knowledge is shaped not just by factors having to do with the world but also by factors having to do with the knower (epistemic freedom) is recognized by many epistemologists. Kant, and, in a different way, contemporary psychologists, linguists, and cognitive scientists, view mental structures as playing a crucial role in determining our theory of the world. The logical positivists emphasize the importance of pragmatic factors in shaping our knowledge. Following Kuhn, many philosophers, historians, and sociologists of science see scientific knowledge as largely determined by social, psychological, political, and environmental factors. The new model accepts the general principle underlying these epistemologies. It follows from this principle that knowledge is not a simple picture of reality (words are not mere labels of objects; sentences and theories are not mirror-images of “facts”), but numerous factors residing in the knower intervene between our world theory and the world.

Critical Note: The existence of factors residing in the knower—epistemic freedom—does not conflict with the applicability of the norms of truth. A street map is not a mirror-image of a city, but relative to certain well-defined goals and standards, it is either a correct or an incorrect image of it. Likewise knowledge, in the present model, is not a *mirror* representation of reality, but it is (or is required to be) a *correct* representation of it. The recognition of this duality is characteristic of NAS. Our statements and theories, according to NAS, are both conventional and factual, both governed by the norms of “simplicity”, “familiarity”, “scope”, and “fecundity” (Quine 1955: 247) and subject to the norms of truth and evidence. This duality extends to ontology: the objects of our theories, as Quine puts it, are both “posits” and “what there is”; as “posits” their knowledge conforms to the norms of utility, as “what there is”—to the norms of truth.⁷ The model explains this two-sidedness of knowledge by its double

⁶ Neither is truth, though truth is correspondence. (See discussion in Chapter 8, Section 3.)

⁷ An especially fruitful use of this duality involves the creation of posits as a tool for studying “what there is”. (See discussion of *composite correspondence* in Chapter 8, Section 4.)

origin in the world and the mind. The interaction of “worldly” and “human” factors is treated not as a mystery by the model, but as something for theories within it (epistemology, psychology, biology, sociology, etc.) to account for.

B. *Multiplicity of Factors Concerning the World (the Factuality of Mathematics, Logic, and Ontology)*. The model admits not only multiple factors residing in the knower, but also multiple factors residing in the world. The multiplicity of worldly factors is related to the diversity of the world. Different parts, features, and aspects of the world are the target of our knowledge, and their diversity partly explains the plurality of our theories. It is commonplace to say that physics investigates a different “region” of reality from biology; but in the new model this view is extended to logic, mathematics, and philosophy as well. Scientists, from Copernicus to Einstein, regarded physical reality as exhibiting not only material features but also *formal* structure, and the factuality and distinctiveness of mathematics are predicated, in part, on the reality and distinctiveness of such structures. The view that logic, too, might be grounded in certain features of reality is suggested by Quine, as we have seen above. In speaking about these features, he says:

[A]dmittedly [logical truth] depends upon none of those features of the world that are reflected in the lexical distinctions; but may it not depend on other features of the world, features that our language reflects in its grammatical constructions rather than its lexicon? (Quine 1970/86: 95).

Ontology, according to Quine, is also concerned with certain features of reality, namely, properties of objects of an especially broad nature:

The question what there is is a shared concern of philosophy and most other non-fiction genres. . . . What distinguishes between the ontological philosopher's concern and [the zoologist's, physicist's, and mathematician's concerns] is . . . breadth of categories (Quine 1960: 275).

Critical Note: In considering the claim that logic and mathematics are (in certain respects) located in the periphery, it is important to note: (i) this claim does not imply that logic and mathematics are empirical; (ii) this claim is not tantamount to Platonism. In fact, neither extreme empiricism (the view that apriori, all knowledge is restricted, or reducible, to empirical knowledge) nor Platonism (the belief in a separate reality of abstract forms) is compatible with our model.⁸

C. *Multiplicity of Factors Underlying Institutions (the Language Institution: Reconciling Holism and Compositionality)*. The principle of multiplicity of factors applies not just to the knower and the world but also to the institutions of knowledge. Take the institution of language, for example. The model permits us to

⁸ We will discuss Platonism further in Chapter 5, Section 2 and Chapter 9, Section 2, and we will further discuss empiricism in Chapter 5, Section 3 and Chapter 6, Section 1. We will offer our own account of the factuality of mathematics and logic in Chapter 5, Section 2, Chapter 8, Section 4, and Chapter 10, Sections 2, 3, 4, 6, and 9.

affirm both the holistic nature of language (the fact that the meaning of any term may, in certain respects and under certain conditions, be affected by that of any other term), and its compositionality (the fact that the meaning of complex terms depends, in some systematic way, on the meaning and mode of composition of simpler terms).⁹ The two are allowed to coexist as two independent factors whose complex pattern of interaction is a matter for linguistics, psychology, sociology, the philosophy of language, and other theories within the model to investigate. By combining holism with discrete constituents the model is “organic” in Dummett’s sense.¹⁰

D. Multiplicity of Factors as Richness of Structure / Enhanced Explanatory Power. The multiplicity of factors characteristic of the new model amounts to richness of structure. The model portrays the process of knowledge as an intricate process, involving numerous factors and constituents dynamically interacting with each other. By allowing rich structures, the model’s ability to deal with complexity and diversity is enhanced. In particular, the model is capable of reconciling elements which, in a simpler model, would appear to conflict with each other—for example, holism and compositionality, as discussed earlier. Indeed, the difference between the old model and the new one can be partly explained as a difference in richness of structure. By taking into account a relatively small number of factors, the old model delineates an overly simple structure of knowledge—a fairly static structure with largely fixed center and periphery, simulating an analytic-synthetic division (albeit in a toned-down and gradual manner). But by increasing the free movement of units of knowledge in two dimensions—context and time—the new model introduces new factors and structure to the old model, relativizing its center-periphery duality to time and context, and thereby reconciling it with NAS. It is important to note, however, that the added complexity does not render the model chaotic. It brings it in line with the complex reality of knowledge, but it does so by introducing further structure rather than unmanageable chaos.

4.2 Temporal-Developmental Dynamic

I will note three types of diachronic change, corresponding to three aspects of the center-periphery duality: (1) change from conceptual orientation to factual orientation and vice versa; (2) movement from “rear” to “front” of knowledge and vice versa, and (3) change from being an invariant (“fixed”) to being a variant (“unfixed”) constituent of knowledge and vice versa.

⁹ My use of “compositionality” here is broader than usual. It covers any semantics which analyzes complex structures by a finite algorithm based on their parts, including, in addition to Tarskian semantics, Henkin’s (1959) functional semantics for partially ordered quantifications, Hintikka’s (1995) game theoretic semantics for IF languages, and my own (Sher 1997b) algorithm for generalized branching quantification.

¹⁰ For another attempt to reconcile compositionality with holism, see Pagin (1997).

I. Change from conceptual to factual orientation (and vice versa)

A. Developmental Change. At each stage in its development, each (viable) science has both factual and conceptual interests, but at different periods in its development different interests may play a more, or less, dominant role. How to classify each period in the history of a given science is a matter for the historian and philosopher of science to determine. But during periods of major conceptual advances the science moves closer to the center; during periods of intensified factual investigations (either in the sense of experiment and observation, or in the sense of theoretical development) it veers toward the periphery.¹¹ Logic and mathematics, too, alternate between the factual and the conceptual: the introduction of the modern quantifiers by Frege, for example, is naturally viewed as a conceptual contribution, while Gödel's discovery of the incompleteness of arithmetic is, in an important sense, a factual contribution. All scientific revolutions, according to Kuhn, involve major changes in conception and methodology (both practical methodology and theoretical methodology); such changes are represented in the model as occurring (largely or partly) in the center.

B. Change in Grounds and Significance. The same statement or theory can be accepted based on one ground during one period, based on another during another (e.g., based on conventional grounds during an earlier period, based on factual grounds later on, as pointed out by Quine in his Genetic Fallacy argument). Similarly, the significance of a statement or a theory may vary in the course of time. For example, it is quite natural to view non-Euclidean geometry as moving from a phase in which its significance was primarily conceptual (independence of the parallel postulate) to a phase in which its role and significance are largely factual (a theory of the curvature of physical space).

II. Movement from "rear" to "front" of knowledge (and vice versa)

A. Movement of Whole Disciplines. In the battle against Nature sometimes one discipline, sometimes another, occupies the front. It is common to think of the natural and social sciences as occupying this position: physics in 1687 and 1905, biology in 1859, psychology in 1900, and so on. Traditionally, logic and philosophy are regarded as permanently located in the center, "setting the stage" for the battle against Nature rather than taking an active part in it. The non-bifurcational methodology of the present model challenges this tradition. Speaking of logic's role in contemporary physics, Finkelstein says:

[L]ogic... is... a dynamic ingredient in the physical theory, **an actor rather than part of the stage** (Finkelstein 1969: 213).

¹¹ Although theoretical and conceptual developments in our model are interconnected, the difference between them is not obliterated. (For a relevant discussion of conceptual advances in science see, e.g., Thagard 1992.)

In the new model logic plays both the role of an actor and that of managing the stage. As an actor, the challenges taken up by logic are of a general and fundamental nature: the challenge of valid inference (transmission of truth from sentences to sentences), the challenge of the basic structure of phenomena (bivalence or non-bivalence), etc. Philosophy, too, is assigned a double role. One of the challenges posed by Nature to philosophy is the skeptical challenge. The skeptical challenge questions humanity's ability to break out of its "mental cage" and attain *true* knowledge of the world. Philosophy took up the skeptical challenge in Descartes's *Discourse on Method* and *Meditations* (1637 and 1641), Kant's *Critique of Pure Reason* (1781/7), and elsewhere. Quine's manifesto, "Epistemology Naturalized" (1969a), calls for fighting the skeptical battle on the psychological rather than the philosophical front. Critics of Quine's naturalism claim the challenge cannot be adequately met by psychology. The new model regards the skeptical challenge as a genuine philosophical challenge, but it does not deny that some parts of this challenge are best dealt with by other sciences. The choice of strategy for dealing with particular skeptical challenges is an open question, to be decided based both on the nature of the particular challenge and on the availability of resources. The balance of these factors may vary from era to era, and with it the choice of strategy may vary.

B. *Movement of Special Issues.* Not only whole disciplines but also single questions can shift their position from the rear to the front. Take the *Entscheidungsproblem*, for example. As posed by Hilbert in 1900, this problem concerned the existence of a finite procedure for determining the solvability of Diophantine equations. But though a problem of considerable interest for mathematics, it was far from being one of the central issues facing our system of knowledge at the turn of the twentieth century. In the 1930s and 40s, however, the *Entscheidungsproblem* became the focus of a serious challenge (or group of challenges) to our system. This challenge—subsumed under the headings of completeness, decidability, computability, recursiveness, etc.—concerns the scope and limits of a broad class of procedures and functions of special importance for knowledge. With the work of Gödel, Church, Turing, Post, von Neumann, and others, the *Entscheidungsproblem* moved to the front of knowledge, leading to some remarkable gains in the battle against Nature, as well as to some serious losses.

III. "Fixing" and "unfixing" constituents of knowledge

The idea of the center and the periphery as representing the "fixed" and "varying" constituents of knowledge, respectively, is a traditional idea: our language, logic, mathematics are *fixed*, experiential knowledge is *variable* and *changing*. The new model relativizes this duality both to context and to time. The underlying principle is Neurath's: we build our system from within by shifting our foothold from one area to another. Placing our foothold in area A—i.e., holding the concepts, laws and theories of area A fixed—we rebuild area B; shifting our foothold to area B (or to an area

intersecting B), we change area C, and so on. The area we hold fixed at time-interval t is the center at t , the area we change at t —the periphery at t . This Archimedean principle is doubly dynamic: we can use A as a lever for changing B both at certain times and in certain contexts. Among the many ramifications of this principle are:

A. *Explanation of Continuity Through Change.* One of the main challenges facing any epistemology is the explanation of continuity through change. The problem is especially difficult for theories permitting radical change, like Quine's epistemology with its universal revisability and inseparability theses, and Kuhn's theory of scientific revolutions as paradigm shifts. In the current literature, the problem is variably referred to as the "incommensurability" problem, the "meaning variance" problem, the "change of theory is change of subject" problem, etc. The new model is well equipped for handling this problem. Think of the "fixed" area of our system at t —the center at t —as containing those elements which are stable at t . The center is what "glues" the system together; the periphery is the locus of change and revision. During periods of "normal" development, logic, methodology, and general scientific principles are held fixed in the center; during periods of radical change, it is conceivable that some combination of simple observation statements, everyday concepts, common-sense maxims, and logical and theoretical elements (not undergoing change) constitute the center. The exact dynamics of variance and invariance is a research project for various disciplines within the model, but the model itself offers a flexible framework for the explanation of continuity through change.¹²

B. *The Foundation Project.* While Quine's original model is not very hospitable to the foundational project,¹³ the new model is. By allowing, and indeed requiring, all disciplines to occupy the periphery (during some periods and in some contexts), the model subjects all disciplines to the norms of *veridical evaluation* and *justification*. And if we regard these as the foundation norms, then every science and discipline within the model—including sciences which in traditional models were placed at the bottom of the foundationalist hierarchy (and hence could not be given a significant foundation)—are now subject to the foundation norms. Furthermore, the holistic structure of the model makes it possible for each science to acquire a foundation. Each science can forge multiple cognitive routes to reality, and each stands in a multitude of normative, conceptual, practical, and theoretical relationships to other sciences. By utilizing these resources rationally yet flexibly we are able to provide it with a significant foundation.¹⁴ While an *apriori*, eternal foundation—a foundation guaranteed in advance for all contexts and times—is at variance with the model, its norms of unity, generality, and economy favor a broad and unified foundation over a narrow or piecemeal foundation. In these ways the model exemplifies the

¹² Similar approaches to the continuity-through-change problem were suggested by Hesse (1970) and by Gochet (1986, Chapter 1, Section 6).

¹³ See Chapter 6, Section 1.

¹⁴ In Chapter 10 we will explain in detail how logic can acquire such a foundation.

foundational-holistic methodology developed in Chapter 2, Section 2, as well as the principles of epistemic friction and freedom underlying this methodology.

C. *A Standpoint for Philosophy* (“*Immanent Transcendence*”). The view that philosophy requires an external standpoint has led such philosophers as Wittgenstein and Rorty to deny the possibility of a substantive philosophy. A substantive theory of knowledge, it is said, requires a viewpoint encompassing our world-theory in its entirety, but such a viewpoint is not available to humanity. The new model enables us to reconcile the immanence of knowledge with the transcendence of philosophy. The model suggests the possibility of a standpoint affording a broad view of knowledge: a standpoint in an area that offers a view of all other areas, both of their relations to each other and of their relation to the world (as seen from the chosen standpoint in the chosen area). In addition, the chosen area may offer an indirect view of some of its own regions, as does arithmetic with respect to its own syntax and certain mathematical and scientific theories with regard to their own ontology (see Quine 1969b). Thus, although an absolute transcendent viewpoint is ruled out by the model, a relative transcendent viewpoint—a point of view transcending all but a makeshift center—is compatible with it.¹⁵ Such a point of view is, indeed, characteristic of the model itself: the model offers a broad outlook on our system of knowledge from a standpoint within it, namely that of a late twentieth-century, early twenty-first-century epistemology.

This concludes the sketch of the dynamic structure of our model. The model exhibits the features (“building blocks”) we were looking for in a model of knowledge: a systematic, yet flexible structure; broad, significant interfaces with both reality and the mind; a rich, holistic network of interconnections, encompassing both connections among units of knowledge and connections between units of knowledge and reality; non-bifurcation of units of knowledge into factual and conventional; a broad, yet substantial, conception of mind, reality, knowledge, and justification; high standards of substantiveness and veridicality for all disciplines; a flexible and dynamic conception of epistemic inquiry; and a significant role for active freedom in the project of knowledge.

In Chapters 5 and 6 we will further elaborate the basic principles of the new model, clarify its differences from Quine, and demonstrate its immunity to prevalent criticisms of Quine’s model.

¹⁵ For further discussion of transcendence see Chapter 8, Section 1.

Reality, Intellect, Realism

One recurrent motif of the present essay is that elements which are naturally viewed as opposed to each other do not conflict after all: epistemic friction does not conflict with epistemic freedom, holism does not conflict with the foundational project, being in the center does not conflict with being in the periphery, being governed by veridical norms does not conflict with being governed by conceptual or pragmatic norms. Our conception of *reality*, *intellect* (*reason*), and *realism* furthers this motif. Our model is characterized by:

1. *A broad and open-ended conception of reality, neither Platonist nor nominalist, one that affirms both experiential and abstract features of reality—one reality—and regards them as interconnected.*
2. *A view of human intellect as playing a central role in knowledge, both abstract and empirical. Intellect's role is central not just to the conceptual or pragmatic aspects of knowledge, but also, and significantly so, to its veridicality. By affirming the central role of intellect in knowledge, however, the model does not affirm apriorism, nor does it identify non-apriorism with empiricism. The model is neither apriorist nor empiricist, regarding intellect and sensory perception as two essential, yet interconnected, elements of knowledge.*
3. *"Basic" realism: A realism which is more robust than most other forms of realism in regarding reality as both the target and ground of all human knowledge (including abstract, e.g., logical and mathematical, knowledge), yet is more flexible than most other forms of realism with respect to the ways in which a theory can be (substantially) connected to reality.*

In going beyond the common dichotomies involving reality, reason, and realism, our model extends a trend begun with Kant's transcendence of the realism-idealism schism, continued with Quine's rejection of the analytic-synthetic dichotomy, and followed by further developments: objections to the apriori-empirical bifurcation (e.g., Railton 2000, Hawthorne 2007, Williamson 2013) and replacement of the traditional apriori by new, mitigated conceptions of the apriori (e.g., Peacocke 2000, Friedman 2001), recognition of subtleties involved in the realism-antirealism debate (e.g., Dummett 1963, Putnam 1987, 1999, Wright 1992, Fine 2001), and so on.¹ But our

¹ By talking about a trend I do not mean to say that all participants in this trend share the same views or that my own views are exactly the same as those of the other participants. Still, I see a trend here.

model has some commonalities with the attitudes and views of those who defend the traditional categories as well: the seriousness with which Platonists approach abstract knowledge and their insistence on the need to ground it in reality; the recognition, by current apriorists (BonJour 1998 and others), of intellect's central role in both the acquisition and veridical grounding of knowledge; the recoil, by nominalists (Azzouni 2004 and others), from ontological redundancies; and the appreciation, by all empiricists, of the crucial role played by the empirical method in knowledge.

In discussing reality, reason, and realism, I will approach these topics in a way that differs in several significant respects from the current approaches. First, many contemporary philosophers focus on relatively small units of knowledge (term, sentence, etc.), or at least on delimited units (specific theories or fields of knowledge). In contrast, I will start with our system of knowledge as a whole. Second, contemporary philosophers often use language, including natural language, as a guide for ontology. In contrast, I regard language as an institution that (to a large extent) evolved relatively early in the history of human culture, was shaped by multiple factors, and fulfills multiple roles; as such it is unlikely to be an adequate guide for ontology, especially where the ontologies of advanced and complicated theories are concerned. Third, in studying the above topics, contemporary philosophers often begin with traditional philosophical categories such as "apriori knowledge" and ask if anything falls under them. In contrast, I will start by looking at the world on the one hand and at our basic cognitive situation on the other, and ask whether it makes sense to focus on the traditional categories in developing our epistemic theories. One advantage of approaching the topics of realism, reason, and reality in this way is that it affords a broader and more basic perspective than those sometime adopted in the current literature.

In pursuing the topics of this chapter I will alternate between an external and an internal point of view. Starting with general observations and considerations, I will incorporate them into my epistemic model/theory as principles, then examine the model (including these principles) from an external perspective, further developing (changing, adjusting) it based on this examination, and so on. In this way, the external perspective will be integrated into the model, becoming internal to it.

For the sake of concise expression I will sometimes speak in my own, first-person voice and sometimes in the third-person voice, treating the model itself as an individual and saying *it* "regards", "holds", "claims" certain things, etc.²

² Someone might wonder whether continuing to develop our model of knowledge in the present chapter is compatible with our view of that model as open to revision. It is important to note that there is no conflict here. In the same way that scientists' belief that their theories are open to revision does not—and should not—prevent them from further developing these theories, defending them against criticisms, clarifying them, and requiring compelling reasons for revising them, so philosophers' belief that their epistemic theories are open to revision should not prevent them from doing any of these things. Universal revisability is not a reason for ceasing to engage in a full-fledged epistemic project. On the contrary, it is only by engaging in this project that the revisions we eventually decide to make amount to progress rather than to capricious, arbitrary, or otherwise unwarranted change.

5.1 The Basic Epistemic Situation

My starting point is the observation that our basic cognitive, and in particular, epistemic situation is characterized by a constellation of factors pulling in different, sometimes opposing, directions. Four such factors, or clusters of factors, are:

FACTOR 1: *The complexity, diversity, and changeability of the world (target of knowledge).*

FACTOR 2: *Humans' ambition to know the world in all (and despite) its complexity, diversity, and changeability.*

FACTOR 3: *The existence of severe physical, biological, psychological, cultural, and other limitations on humans' cognitive capacities.*

FACTOR 4: *The richness, power, intricacy, flexibility, malleability, maneuverability, and continual growth of humans' mental resources (concurrent with Factor 3).*

The combination of these factors affects both the scope and the limits of human knowledge, both its challenges and its promise. It further affects our choice of methods for pursuing knowledge. The crux of the matter is that we humans seek to know things about the world which might not be easily, directly, or naturally³ accessible to us, and this has a significant impact on the way we go about developing our theories. The challenge we face can be described by saying that if the world is highly complex relative to our cognitive capacities, and if we nevertheless seek to know it in its full complexity, then we must stretch our cognitive endowments, devise multiple means for reaching those regions of the world which are less accessible to us, improvise, experiment, exercise our imagination, etc., in order to attain (or maximize the attainment of) our goal. In short, it requires us to devise, and pursue, a wide variety of cognitive routes from mind to world, some of which are going to be indirect, complex, intricate, jagged. The same thing applies to our use of language, a tool that, for the most part, was bequeathed to us by our forebears, and over which we do not have full control. The most effective use of, say, a *singular term* in a complicated cognitive endeavor might not be to target an *individual* in the world, but rather to target something else in the world, something that, for one reason or another, is more fruitfully studied when represented as an individual and referred to by a singular term.⁴ The key idea is that our theories are created *by and for humans*, and as such might convey information about the world in ways that are advantageous for humans but circuitous from the point of view of a given subject matter, considered

³ I use "natural" here in the sense in which it is, say, more natural for us to think of space in Euclidean terms than in Riemannian terms.

⁴ Such a singular term might target a property of individuals, a property of properties of individuals, etc., and the reason for using it in this way might have to do with our cognitive makeup, history, and so on. We will return to this point in the next section, as well as in the discussion of mathematical truth in Chapter 8, Section 4, where a detailed example will be given.

on its own. Being circuitous, however, does not mean being sloppy, piecemeal, or amateurish. On the contrary, circuitous theories can be as *accurate and systematic* as other, more straightforward theories, and the license to construct such theories increases the likelihood of bringing those aspects of reality which are not directly accessible to us into the circle of human knowledge.

Looking at our own discipline, we observe that epistemology itself is a field of knowledge, directed at a certain aspect of the world, namely, the pursuit of knowledge by the world's human inhabitants. As such, epistemology investigates the human epistemic condition. Yet to investigate this condition effectively, we, epistemologists, must take this condition itself into account while designing our investigations. This is in principle possible if we use the foundational-holistic methodology, but it requires that we give careful thought to the question of what concepts, categories, and dichotomies we use in our investigation. While some concepts, categories, and dichotomies might contribute to the investigation, others will stand in the way. For example, if neither sensory perception *by itself* nor intellectual cognition *by itself* plays a dominant role in generating, explaining, or justifying most types of knowledge, if only certain *combinations* of these two sources successfully perform these tasks, it will be unfruitful to limit the resources used in our investigation either to purely sensory resources or to purely intellectual resources. And this means that neither the concept of the *purely intellectual* (or the *apriori*) nor the concept of the *purely sensory* should guide us in designing our investigations.

These observations on our basic epistemic situation and its challenges are not intended to replace the careful study of the human epistemic condition by various disciplines, from neuroscience to epistemology. Rather, their purpose is to provide a general orientation for our epistemic investigations, a compass which will guide us in our study. In a sense, these observations have guided me throughout this work, and are thus integrated into the dynamic model and foundational-holistic methodology developed in earlier chapters. In seeking to fathom reality, we must use a wide variety of mental capacities. But although these capacities are not of our making, they are not completely beyond our control. The development of new areas of knowledge, new theories, and new ideas is often accompanied by new uses of our capacities: the creation of new concepts, the generation of new methods of justification (proof, experimentation), the use of new perspectives, etc.; in short, the development of new cognitive tools of a variety of kinds. All this means that the route from thought to reality is intricate, complex, and multifaceted, a route that takes multiple forms and is constantly evolving, i.e., a holistic and dynamic route (in our sense of these attributes).

The grounding of our model in humans' basic epistemic situation explains why the model, though primarily normative, is also to a certain degree descriptive. Since in pursuing knowledge humans have to face the problems generated by their basic situation, what they do when they actually succeed in this task often coincides with what our model says they should do.

5.2 Reality: Beyond Platonism and Nominalism

One of the questions that stand at the center of the philosophical literature on the scope and nature of reality is whether there are abstract (as contrasted with physical) objects in the world, or more generally, whether reality has abstract features. The consequences of an answer to this question for the scope of knowledge are potentially considerable, especially on the view that knowledge qua knowledge is knowledge of reality or of some aspect of reality. If reality does not have abstract features or aspects, then whole disciplines, including mathematics and logic, are either not fields of knowledge at all, or are fields of knowledge of an aspect of reality which differs radically from the one they purport to study.

Two prominent doctrines that offer conflicting answers to this question are Platonism and nominalism. While the former affirms the existence, in principle, of abstract objects, the latter denies it. A comprehensive discussion of these doctrines is beyond the scope of this work, but briefly, we can describe them in a way that is pertinent to our concerns as follows:

A. *Nominalism*. Contemporary nominalists typically affirm the existence of physical objects and deny the existence of all other kinds of objects. In this sense their nominalism is associated with extreme empiricism. The motivations for such nominalism can be divided into five clusters:

- (a) Belief in physicalism, priority of science over mathematics, causal theory of knowledge, and naturalistic conception of epistemology.
- (b) The principles of economy, simplicity, and Occam's Razor ("don't multiply entities beyond necessity").
- (c) Preference for a "hard-headed, no-nonsense" (Burgess 1983: 93), common-sensical style of philosophy.
- (d) Problems with alternative theories, in particular Platonism.
- (e) The allegedly mysterious or unintelligible nature of abstract objects. Burgess and Rosen (1997: 29) point to the following views as characteristic of nominalism:
 - (i) "[R]eality is... a system connected by causal relations and ordered by causal laws".
 - (ii) There is no reality outside space and time.
 - (iii) Platonism is "an especially unattractive... supernaturalism".
 - (iv) A "great gulf or... wall" separates human knowers from the Platonic reality. This wall is "causally unpenetrable", hence "flesh-and-blood subjects" are incapable of attaining knowledge of what there is "on the other side of" the wall.

A typical argument in support of nominalism is delineated by Burgess:

All entities of which we can have knowledge are causally connected with our organism.

No abstract entities are causally connected with our organism.

Ergo: No abstract entities are entities of which we can have knowledge (Burgess 1983: 100).⁵

Among the problems facing nominalism, the following are especially pertinent to our interests:

- (a) Nominalism overlooks important philosophical considerations, most importantly, considerations of *truth*. This is reflected in the fact that nominalists rarely ask questions like: “Is it true that reality has no abstract features?”; “Is it true that humans have no resources for gaining knowledge of such features?”. Positive answers to these questions have never been established, yet without establishing such answers it is hard to justify nominalism’s viability.
- (b) Nominalism’s strong suits—common sense and economy—are more limited than they appear to be. Common sense, as we will shortly see, strongly suggests that reality does have abstract features and that humans do have knowledge of such features. And while ontologically, nominalistic reconstructions of physical and mathematical theories are more economical than standard versions of these theories, overall they are often less economical.
- (c) Nominalists’ conception of human knowledge is very narrow. This pertains, among other things, to their insistence that causal connections are the *only* source of human knowledge and to their obliviousness to intellect’s *central role* in knowledge. (Nominalists acknowledge intellect’s role in the conceptual and pragmatic components of knowledge, but not in factual discovery or in factual justification). As noted in (a) above, nominalists rarely offer a thorough justification of their narrow view of knowledge.⁶

B. *Platonism*.⁷ Contemporary Platonism focuses on *abstract objects*⁸. A paradigm of abstract objects is mathematical individuals, like the natural numbers. In a recent encyclopedia article Linnebo describes mathematical Platonism as follows:

Platonism about mathematics (or *mathematical platonism*) is the metaphysical view that there are abstract mathematical objects whose existence is independent of us and our language, thought, and practices. Just as electrons and planets exist independently of us, so do numbers and sets. And just as statements about electrons and planets are made true or false by the objects with which they are concerned and these objects’ perfectly objective properties, so are statements about numbers and sets (Linnebo 2009/13: 1).

⁵ The display of the argument—each premise on a separate line, etc.—is different from the original.

⁶ For the claim that nominalists fail to attend to the justification of their conception of knowledge see Burgess (1983) and Burgess and Rosen (1997).

⁷ Some contemporary authors use “Platonism” for Plato’s original theory and “platonism” for contemporary theories. I will use “Platonism” for the latter, which are the only theories I consider.

⁸ Whereas earlier versions of Platonism often focused on *universals*.

Two motivations for Platonism have to do with:

- (i) the view that many mathematical and other abstract statements (theories) are true, and
- (ii) the view that language is a reliable guide for ontology.

In recent literature these views are often combined. Thus, Linnebo says:

The most important argument for the existence of abstract mathematical objects derives from Gottlob Frege and goes as follows (Frege 1884). The language of mathematics purports to refer to and quantify over abstract mathematical objects. And a great number of mathematical theorems are true. But a sentence cannot be true unless its sub-expressions succeed in doing what they purport to do. So there exist abstract mathematical objects that these expressions refer to and quantify over (Linnebo: 1).

And Balaguer (2004/9) describes the “singular term argument” for Platonism as follows:

The general argument strategy here has roots in the work of Plato, but its first clear formulation was given by Frege (1884, 1892, 1893..., and 1919)... [A] general formulation of the argument [is]:

- 1. If a simple sentence (i.e., a sentence of the form ‘ a is F ’, or ‘ a is R -related to b ’, or ...) is literally true, then the objects that its singular terms denote exist. (Likewise, if an existential sentence is literally true, then there exist objects of the relevant kinds....)
- 2. There are literally true simple sentences containing singular terms that refer to things that could only be abstract objects. (Likewise, there are literally true existential statements whose existential quantifiers range over things that could only be abstract objects.) Therefore,
- 3. Abstract objects exist (Balaguer 2004/9: 14).

Now, as noted in the opening section of this chapter, there is something about Platonism that we full-heartedly endorse: namely, its serious attitude toward abstract knowledge. This includes its view of such knowledge as genuinely true and its view of abstract features of objects as real. But a distinctive characteristic of Platonist theories is the deep gulf they introduce between the physical (or the concrete) and the abstract. One line of reasoning leading to this gulf is the following:

- (a) Abstract objects, being abstract, must exist entirely outside of space and time.

Therefore:

- (b) They must exist in a reality different from the reality in which physical objects exist.

Hence:

- (c) To account for their existence, we must affirm a second reality, a Platonic or abstract reality, of a completely different kind from physical reality.

Although Platonists' talk of "two realities" is often metaphorical, their idea of a radical division between the abstract and the "concrete" is not.

From our perspective, Platonism faces at least five problems: (i) the problem of frictionless theorizing, (ii) the bifurcation problem, (iii) the cognitive access problem, (iv) the application problem, and (v) the language-as-an-arbiter-of-ontology problem. To sharpen the formulation of these problems, let us relate to an extreme form of Platonism, one that advances a strong "two realities" ("two worlds") thesis. For such a Platonism, we can describe these problems briefly as follows:

(i) *Frictionless theorizing.* As we recall from Chapter 1, Section 1, Kant viewed traditional Platonism as a paradigm of a frictionless theory. But all Platonist theories face the challenge of frictionless theorizing to a significant extent. It is too easy to solve philosophical problems by postulating a new reality, without subjecting the resulting theory to the same constraints, or analogous constraints of the same magnitude, as those we set on other theories.

(ii) *The bifurcation problem.* Platonism is a bifurcational doctrine. By postulating two separate realities, a physical reality and an abstract Platonic reality, it gives rise to a bifurcated metaphysics which, in turn, induces a bifurcated epistemology. That is, assuming, as we do, that knowledge qua knowledge is knowledge of reality, Platonism introduces a deep gulf between physical and abstract knowledge. In Chapter 3 we pointed out some of the problems besetting the empiricist bifurcation of our system of knowledge. Let us now proceed to two general problems that are exacerbated by the Platonist bifurcation.

(iii) *The cognitive access problem.* The problem of cognitive access is a general epistemological problem that has at least two dimensions. First, there is the issue of how researchers can cognitively reach abstract features of reality, and second, there is the issue of how epistemologists can provide a grounding in reality for theories of those features. Platonism considerably increases the level of difficulty of this problem. Let me explain: if the world we live in has both concrete and abstract features, then it will not be surprising if we, who evolved, developed, and exist in this world, are endowed with cognitive capacities for accessing, and developing theories of, both these features, and it will not be too difficult to see how we can ground these theories in the world. But if we lived in a world that had only concrete features, then it would be quite surprising if we were endowed with cognitive capacities for accessing abstract features of objects, features that no objects in our world had, but that objects in another, altogether different and separate world, had, and it would be very difficult to see how we could ground any of our theories in that world.

(iv) *The application problem.* Here, too, the difficulty is to explain how laws that hold in one Platonic reality are applicable to another separate physical reality. Furthermore, if the two realities are radically different, why should we surmise that laws governing one reality are instantiatable in the other reality?

Now, of course, if one construes Platonism in a way that does not involve two worlds and does not introduce a sharp bifurcation between abstract and concrete knowledge, then these special difficulties will not arise. Whether, and to what extent, contemporary Platonist theories are committed to such a bifurcation is an issue I will not attempt to adjudicate here.⁹ The treatment of logical and mathematical knowledge in this work affirms their abstractness without requiring extreme bifurcation.¹⁰

The last problem facing Platonism, from our perspective, is:

(v) *The language-as-an-arbiter-of-ontology problem.* The assumption that language is an arbiter of ontology plays a significant role in motivating Platonism: it is partly due to this assumption that philosophers feel compelled to treat the referents of, say, numerals as individuals. But while it is reasonable to expect that objects in our world have, in addition to concrete features, also abstract features, it is difficult to justify the assumption that the world we live in contains not just objects with abstract features, but objects which are entirely abstract and, in particular, abstract individuals, like numbers. A seemingly natural solution is to place numerical individuals in a different world: numerical individuals do not exist in our world, but they do exist in a different Platonic world.

This move is natural *if* we assume that language is an arbiter of ontology. But is it reasonable to assume it is? Given the complexity of the world relative to our cognitive capacities on the one hand and the multiple factors, needs, uses, and mere accidents that shaped human language on the other, it seems to me unreasonable to presume that natural language, and especially its rigid grammar, are automatically suitable, as they stand, for rigorous theorizing in every field. In particular, it is unlikely that a highly specialized and late-coming human activity like rigorous theory-building always abides by a simple 1–1 correspondence between syntactic and ontological types. This is supported by the observation that while in recent centuries the pace of change in our understanding of the world (including what kind of things exist in the world) has been quite fast, the pace of change in our vocabulary has been slower, and the rate of change in our grammar or syntax has been extremely slow. Indeed, it is not clear that such a change has occurred at all, at least in the natural or quasi-natural languages in which most theories are actually formulated. It is to be expected, therefore, that natural language—and in particular its syntax—is at present *not* a very good guide to ontology, and that in figuring out how a given contemporary theory is anchored in reality we have to be flexible in considering the space of possible connections between its linguistic formulation and the world.

Furthermore, it is not even clear that the artificial languages of logic are so built as to provide an accurate regimentation of *ontology*. After all, these languages are

⁹ Today, many Platonists characterize their position in more moderate terms. For example, instead of talking about *two realities* they may talk about *two realms* or *two aspects* of reality (see e.g., Balaguer 1998: 8). In those cases it depends on the particular characterization whether it is subject to our criticisms or not.

¹⁰ See treatment of mathematics and logic in Chapter 8, Section 4 and Chapter 10.

designed to identify the *logical*—not the ontological—consequences of theories.¹¹ It is thus important to realize that there need not be a direct correlation between the degree of realism of a given theory—how robustly it is anchored in reality—and the degree to which its syntax reflects its ontology.

The possibility of a non-straightforward relation between language and ontology is further supported by the central role of epistemic freedom in human knowledge. Given our cognitive plasticity, our ability to find and create new cognitive routes from thought to reality (including relatively complex routes), our facility in maneuvering between multiple levels of thought, our practice of developing sophisticated cognitive strategies, our continuous deliberations and decision-making, and so on, it stands to reason that these capacities and practices will be exercised in our use of language as well, enabling us to circumvent constraints that syntax and grammar are often thought to impose upon us.¹²

C. The Present Orientation Toward Reality. To develop my own orientation toward reality I will start with a basic question: “Does it makes sense to presume that reality—the reality in which we, human beings, and our environment reside—has abstract features?”. Compare this question to the analogous question about *physical* features. There is a sense in which it is reasonable to presume that reality has physical features, that the reality of these features is the starting point of physical science, and that this starting point is still basically valid for contemporary science. Does the same hold for abstract features and abstract science, e.g., mathematics? Note that I am not asking whether there are *abstract objects* in the world, in a narrow sense of “object”, whose paradigm is “an individual”. This would place me squarely at the center of the nominalist-Platonist controversy. Instead I am asking, in very general terms, whether reality has *abstract features*, where “feature” is used as an everyday notion, one which is more similar to “property” than to “individual”, and even then to “property” in a non-committal sense, a sense that distances us from controversial conceptions of properties, like the traditional conception of properties as *universals*¹³.

Let us then turn to our basic question: “Are there strong common-sensical reasons for presuming that reality—including everyday physical reality—has abstract features?”. What we are looking for, at this point, are a few initial yet compelling

¹¹ For further discussion of this point see Sher (2000, 2016b). Quine’s criterion of ontological commitment will be discussed in this chapter at Section 5.4.

¹² Note that whereas in Chapters 3 and 4 I have emphasized the *connection* between theoretical and linguistic change and the *malleability* of language, here I emphasize the *disconnect* between them and the *rigidity* of language (especially syntax). There is, however, no conflict between the two. While some forces push in one direction, others push in the opposite direction, and whereas in some respects and to some extent language is malleable, in others it is not. What this means is that in understanding theoretical change and its connection to linguistic change we have to take into account both forces.

¹³ Universals are often treated as a kind of *entity*, and as such are embroiled in the controversy between nominalists and Platonists.

observations, ones that will provide us with a fruitful starting point, give us a basic orientation, and introduce a fresh perspective into our inquiries.

Our first observation is straightforward: objects in the world have features of a variety of kinds. The question is whether they have abstract features. We will focus on “formal” features as a paradigm of abstract features, but a precise characterization of “formal” is not needed here.¹⁴ Let us begin with non-controversial objects, objects whose reality is usually not questioned either by nominalists or by Platonists: humans. Consider the students in one of my graduate seminars, say, Truth in Kant (UCSD, 2010, seventeen students). First we observe that each individual student in the class has many features, some abstract, some not. For example, each individual student has the abstract (formal) property of being identical to himself/herself and the non-abstract property of being a student in the class. Next we observe that properties of these students, too, have many features, and some of them are abstract as well. For example, the property of being a student in the class has a cardinality property, SEVENTEEN, which is a formal, abstract property. Proceeding to binary relations of students in the class, we observe that these, too, have some formal features. For example, the relation of studying-in-the-same-class-as has the formal features of being reflexive, symmetric, and non-transitive. Similarly, properties of students in the class can be operated on, or be combined, by various formal operations: complement, union, intersection, Cartesian product, etc. For example, by applying the formal operation of intersection to the properties of being a first-year student and being a woman we obtain the intersective property of being a first-year female student.

We do not need more than plain common sense to make these observations, and it is quite hard to contest them. Suppose that neither my students nor their properties have any abstract features. What does this mean? It means that these students are neither identical to themselves nor different from any other individuals, that collections of these students have no definite cardinalities, that properties of these students do not form unions and intersections, that relations between them exhibit no formal patterns (e.g., no relations between them are reflexive, symmetric, or transitive), and so on. These claims, however, are unreasonable. So if my students are real, and if they have the properties and relations mentioned above, and if these properties and relations have the features noted above, then objects in the world do have abstract features, and such features are real. In the absence of any compelling challenge to the

¹⁴ By starting with the general and flexible category of “formal feature” we gain the advantages of not having to commit ourselves to the existence of abstract individuals (which are the focus of many objections to abstract knowledge and ontology), being able to leave the precise nature of properties to a future work, and being able to arrive at a characterization of formality in a gradual, step-by-step, manner. Thus, below we will give a few intuitive examples of formal features; in Chapter 8, Section 4 we will see how an account of mathematical truth that avoids commitment to formal individuals can avoid influential arguments against mathematical realism; and in Chapter 10 we will offer a precise characterization of formality, show that no individuals fall under this category, and explain its centrality to logic.

existence of my students, to their having such properties as self-identity, and to their properties having such properties as cardinality, the above considerations suggest the reality of abstract features.

Our answer to the basic question posed above is, thus, positive. Yes, it is reasonable to surmise that reality, the reality of regular physical objects like people and trees and mountains, has abstract features. Indeed, it is as reasonable to surmise this as to surmise that they have physical features. This leads to a view that, like Platonism, affirms the reality of abstract features, and like nominalism, has no need for a “second” reality. The reason for the latter is that the abstract features we have observed are features of objects and properties residing right here, in our world, and they hold of these objects (properties) right here as well, rather than in some other reality. That is, the world of bona fide physical features is also a world of abstract features. This view avoids the excesses of both nominalism and extreme Platonism, charting an approach to reality that goes beyond the existent alternatives.¹⁵

5.3 Intellect: Beyond Apriorism and Empiricism

Knowledge of the world, or of various aspects of the world, is attained, in our model, by a combination of activities, ranging from sensory perception to conceptualization, abstraction, generalization, reflection, combinatorics, analysis, figuring out, model building (scientific and logical), experiment design, mathematical intuition, and so on. In this essay I place these activities, with the exception of sensory perception, under “intellect”. Intellectual activity, in the sense intended here, is thus any non-sensory cognitive activity, i.e., any cognitive activity that does not principally rely on any of our five senses. I also use “reason” as a synonym of “intellect” in this sense.

Our distinction between *sensory perception* and *intellect* (in the above sense) is clearly *non-Kantian*, since intellect (in our sense) includes *pure intuition* (in Kant’s sense), but Kant regards pure intuition as closer to sensory perception than to *reason* or *understanding*, which are what he might identify as intellect.¹⁶ However, our distinction is *not anti-Kantian* either, since it is central for us, as it is for Kant, that sensory perception—perception centrally involving the five sensory modes—has a very limited, if crucial, role in human knowledge, and it is only in combination with other cognitive resources (falling under our “intellect”) that knowledge, and even mere cognition, is possible.¹⁷ Furthermore, our demarcation of the resources of

¹⁵ As readers might recognize, there is an Aristotelian element in the present approach. I will discuss both the similarities and the differences between this approach and Aristotle’s in Chapter 8, Section 4. In that section I will also propose an account of mathematics that affirms its abstractness without being Platonist. In Chapter 10, I will do the same for logic.

¹⁶ Kantian understanding also falls under our notion of intellect.

¹⁷ I use “cognition” as a weaker notion than “knowledge”, one that does not involve a claim to justification and possibly does not involve a claim to truth as well. (These are possible interpretations of Kant’s “*Erkenntnis*”, though nothing that I will say here depends on Kant’s notion.) For the purpose of the

knowledge into sensory perception and intellect is intended only as an initial demarcation, not as a final result. What I am seeking (here as in the last section) is a fruitful starting point, one that will free us from the specific commitments of existent theories. This is also the reason I use “intellect” more often than “reason” in this essay. In contemporary philosophy, “reason” is (or at least appears to me to be) a more loaded notion than “intellect”; hence it is easier to keep a distance from its specialized uses by talking about *intellect*.

There is no question of a precise definition of “intellect” here. For one thing, this notion is too broad and complex for a single, simple definition; for another, much research remains to be done in order to understand the workings of intellect in knowledge. The message I would like to convey is that (prima facie) knowledge of the world requires some kinds of rational access to reality along with sensory access.

While the form of intellectual access to reality that has received the most extensive attention in contemporary philosophy is that of mathematical intuition (see e.g., Parsons 1979, 1995), what I would like to emphasize here is a different cognitive activity (or a cluster of such activities), one that may be naturally called “figuring out”, as in “figuring out how things are in the world, how things in fact are”, “figuring out how to manipulate the world”, “figuring out how to justify claims about the world”, and so on.¹⁸ Figuring out in this sense includes both figuring out *that* and figuring out *how*, which are often related.¹⁹ The activities I have in mind as falling under “figuring out” are, thus, primarily activities centered on the world (facts). They are also activities that humans constantly engage in in all areas of life, both practical and theoretical, and, indeed, at all ages, from early childhood on. For example, when, “putting two and two together”, a baby realizes he can make the mobile on his crib move by banging his feet on the mattress as hard as he can, he is engaged in *figuring out*. When, upon being told that a recent incident of sudden darkness at home (which she found thrilling) was caused by using too many electrical appliances at once, a toddler starts turning on one electric appliance after another, she is engaged in *figuring out*. When a technician tries to fix a broken computer, what he does first is *figure out* what the problem is. What a farmer, a builder, an electrician spends much of her time on, day in and day out, is *figuring out* the roots of a problem, how to fix it, how to build new tools (or what tools to buy), how to improve her enterprise (farm, business), and so on. What logicians and mathematicians spend much of their time on is *figuring out* how to solve mathematical problems and how to prove their results. Thus, what Russell did upon discovering a paradox in Frege’s logic was *figure out* how to block it without losing the power of Frege’s logic. What Gödel did in arriving

present section there is no need to sharply distinguish between cognition and knowledge, and I will freely switch from one to the other.

¹⁸ I should add, however, that some of those who focused on intuition introduced elements that are close to my conception. See, in particular, Parsons (2000, 2008) and Posy (2000).

¹⁹ Stanley (2011) suggests one perspective on their relatedness.

at his incompleteness theorems involved a lot of *figuring out*, from *whether* arithmetic was complete to *how* to construct a rigorous proof of its incompleteness (e.g., how to use syntax to represent itself). What Wiles and his colleagues did was to *figure out* that Fermat's last theorem was related to a conjecture originally raised independently of it, the Taniyama-Shimura conjecture, and then *figure out* how to prove the relevant part of that conjecture. What, according to *The Double Helix* (1968), Crick and Watson did when they set out to build a model of DNA was to engage in a series of steps of *figuring out* what, given their present situation, they had to do next in order to approach their goal.²⁰ Einstein's thought experiments were acts of *figuring out*. The design of scientific experiments requires much *figuring out*. *Figuring out* is what contemporary physicists are doing when they try to reconcile quantum mechanics and relativity theory. It is what underlies Kant's Copernican revolution. (Kant *figured out* that one way to establish the possibility of scientific knowledge in the face of Hume's challenge was to change the traditional epistemic coordinates, or make a shift in the hitherto-universal epistemic *gestalt*.) More generally, *figuring out* is what most philosophers, and especially those oriented toward problem-solving, are constantly doing (or are trying to do). And it is, of course, what we are trying to do in this essay: in seeking an alternative to foundationalism that does not renounce the foundational project, in seeking a model of knowledge that shares the advantages of Quine's model while avoiding its limitations, in seeking to understand the nature of truth, in seeking a philosophical foundation for logic.²¹

Some of the activities mentioned in our examples might turn out to be *non-intellectual* activities, or to lack a *significant* intellectual component. This is to be expected, since the above observations are intended to serve as *initial* observations, as something that motivates us in our investigations, rather than as conclusions of completed investigations. But what is already informative about these observations, what comes out of these observations, is (i) that the act of *figuring out* is different from the act of exercising our sense organs, and hence it is unlikely to be explained in purely experiential terms, (ii) that while *figuring out* is not itself a sensory activity, it might appeal to, or work in tandem with, sensory activities, and (iii) that *figuring out* takes place in a large array of contexts, both abstract theoretical contexts and concrete practical contexts. More generally, a proper act of *figuring out* may employ prior knowledge, may have recourse to experience, may involve a variety of mental faculties, and so on. *Figuring out* is thus best understood in a holistic setting, in a setting that sanctions the joint use of multiple cognitive resources to achieve an epistemic goal. Finally, while the above observations are partly psychological, they are largely epistemological. They are observations of, and conjectures about, *what it*

²⁰ In our discussion of epistemic freedom in Chapter 1, Section 3 we described this as a series of decisions.

²¹ Something like "figuring out" can be found in Dewey (1938), Brandom (1994a, 2000), and others.

takes to attain our epistemic goals, what roads are *open to humans* in seeking to attain these goals, both in the practical and in the theoretical arena.²²

But there is a problem with some of the traditional categories often used in philosophical discussions of intellect (reason). Categories like *apriority* are conducive to an excessively narrow view of intellect as a source of knowledge. Such a view focuses on intellect as it operates in *isolation* from other sources of knowledge (in particular, sensory perception) rather than in cooperation with such sources. I do not mean to say that this category, or the associated apriori-aposteriori distinction, are unintelligible, or undefinable, or empty, or not useful for any purpose. But it seems to me that placing this category (distinction) at the epicenter is counterproductive for understanding the role of intellect in knowledge, in particular its role in *factual discovery* and *veridical justification*, the two aspects of knowledge we are interested in here.

To identify the root of the problem I will use a figurative notion, “the line of knowledge”, or more precisely “the line of *cognitive resources* of knowledge”, where knowledge is conceived, as it is throughout this essay, as knowledge of the world, and the cognitive resources have elements of two kinds, *sensory* and *intellectual* (in our sense). Traditional conceptions involve various assumptions about the line. Often, they assume that the line has two end sections or points (see Figure 5.1), representing two significant types of knowledge or cognition, namely, *purely* sensory cognition (A below) and *purely* intellectual cognition (E below).

For the sake of analysis, let us visualize the line as containing 5 sections, A–E, as follows:




Figure 5.1

where

- A = Segment of purely sensory cognition, i.e., cognition in which intellect plays no significant role whatsoever,
- B = Segment of largely sensory cognition, i.e., segment in which sensory resources play the major role, but intellectual resources also play some role,
- C = Segment of roughly equal role for sensory perception and intellect in cognition,

²² Our discussion of figuring out raises the question “What cognitive capacities enable us to figure things out?” I will not investigate this question in this essay. Here my goal is to point out that the category of “intellect” is far broader than it appears to many epistemologists, that it is not limited to purely apriori knowledge, and that some intellectual activities are universal in the sense that they are essential for knowledge in all fields, from experimental physics to logic and mathematics, as well as for many of our everyday and professional activities. I do hope to study intellect in general, and figuring out in particular, in the projected volume *Epistemic Freedom*. More importantly, I hope to kindle others’ interest in such a study which, I believe, is crucial for a better understanding of human knowledge.

- D = Segment of largely intellectual cognition, i.e., segment in which intellectual resources play the major epistemic role, but sensory resources are *not* excluded,
 E = Segment of purely intellectual cognition, cognition in which not even the slightest amount of sensory experience plays an epistemically significant role,²³

and the size of each section represents the number of possible combinations of sensory and intellectual elements that are represented by it. It does not matter here whether B, C, and D are equal in size or not, or if there are sharp boundaries between them, or what exactly the intellectual and sensory elements are. It is, however, significant that A and E are sections that allow only one combination of intellectual and sensory elements each: only sensory elements—A, only intellectual elements—E. In this sense we may view A and E as *points*.

Now, the traditional division between the *apriori* and the *aposteriori* cuts the line of knowledge into two at *one of its ends*.²⁴ The *apriori* segment consists of the single point E, whereas the *aposteriori* segment consists of almost the entire line, from A to the outer reaches of D (its closest points to E). All points on the line but one fall under *aposteriori cognition*, which is traditionally identified as *empirical* cognition; only one point on the whole line falls under *apriori cognition*. If, idealizing, we speak in terms of *ratios* of perceptual to intellectual sources in our knowledge of the world, we can say that the *apriori*-*aposteriori* distinction divides the line of knowledge into two parts. One part consists of knowledge whose sources belong to a very small interval of ratios, namely the singleton interval [0:100], and the other part consists of an interval of ratios spanning almost the entire line, [100:0–1:99]. In this way, the *apriori*-*aposteriori* distinction focuses our attention on one mode in which intellect might work, namely, that in which it works in *complete isolation* from all other sources, and away from the multiplicity and diversity of modes in which it operates in *cooperation* with other sources.

Of course, the distinction itself does not forbid us to study intellect's operations in regions A–D, and there might be some purposes for which focusing on E-knowledge would be useful. But it is not clear, in advance of further investigation, that the A and E “regions” are not empty (that they are not just *ideal limits* of B and D), and in any case it is unreasonable to decide in advance—that is, in advance to serious inquiry—that *apriority* can be used as a *guide for understanding the overall role of intellect in human knowledge*.

To further clarify this point, let me note that there is no connection in principle between intellect playing a *central role in all areas of knowledge* and its being *the only player in some areas of knowledge*. Thus, while we accept Kant's point that:

though all our knowledge begins [in the order of time] with experience, it does not follow that it all arises out of experience,

²³ As before, what I am interested in here is its role in factual discovery and/or veridical justification.

²⁴ A somewhat similar point is made by Williamson (2013: 291–2) with respect to the top-down conception of the *apriori*-*aposteriori* distinction.

we do not see it as implying that:

there is ... knowledge that is ... [completely] independent of experience, ... entitled *a priori* (Kant 1781/7: B1–2).

That is, the claim that intellect plays a central role in knowledge leaves it an open question whether this role is exhausted by, or even significantly involves, apriori knowledge, i.e., whether region E of the line of knowledge is not empty, and if it is not empty, how much significant knowledge is located in it.²⁵

One possible reason for preserving the category of *pure apriori knowledge* is its ability to explain another attribute of some units of knowledge, which we might not wish to discard, namely, *necessity* or *modal force*. But as Kripke (1972/80) showed, the two are not inextricably connected. Furthermore, as we will see in Chapter 10, Section 5, we do not need apriority to account for the strong modal force of those items of knowledge which have such force; there are other, better ways to explain their modal force. For example, we can explain the strong modal force of logic and mathematics based on their *strong degree of invariance*.²⁶ One advantage of using *invariance* to explain modal force is that invariance comes in degrees. Apriority induces a binary account of modal force, one that focuses on two categories, *the necessary* and *the contingent*, while invariance focuses on types and degrees of necessity, which is what we need to understand differences in modal force between different branches of knowledge (due to differences in the modal force of their principles or laws). Thus, invariance explains why physical laws have a weaker modal force than logical laws yet still have a significant modal force. It explains why biological principles have a weaker modal force than physical principles yet stronger than that of observation statements, and so on. The appeal to invariance also explains why logic and mathematics *appear to be* purely apriori and why intellect plays a dominant (though not necessarily exclusive) role in logical and mathematical knowledge. We will say that invariance explains why logical and mathematical knowledge is “quasi-apriori”.

It is important to indicate, however, that by classifying logical and mathematical knowledge as *quasi-apriori* rather than as purely apriori—i.e., as dominantly intellectual but not as exclusively intellectual—we are not denying the central role of activities that are traditionally viewed as purely apriori, like *mathematical proof*, in knowledge. On the contrary, we regard mathematical proof as one paradigmatic example of intellect’s ability to produce *epistemic friction*, and we regard knowledge based on mathematical proof as one paradigm of human knowledge. Where *quasi-apriority* finds its place in such knowledge is not in proof itself, but in *establishing the*

²⁵ Concerning Kant: although Kant himself does not explicitly say that the first point implies the second, it is not clear that he does not implicitly connect the two.

²⁶ The notion of invariance will be explained and discussed at length in Chapter 10, Sections 4–7. This applies to all the comments below concerning invariance (including the one on *quasi-apriority*).

validity of the proof method, in choosing mathematical axioms and logical rules of proof, in grounding mathematics and logic in reality, and so on.²⁷ Nor do we deny that mathematics and logic are grounded largely in abstract features of reality. What we question is the view that there is a complete disconnect between abstract and physical features of reality and, based on it, between logico-mathematical branches of knowledge and physical branches.

Our emphasis on the centrality of intellect to knowledge on the one hand and the shortcomings of traditional apriorism on the other connect us to views expressed by other contemporary philosophers on these issues. These include, as we have pointed out above, Railton (2000), Hawthorne (2007), and Williamson (2013), among others. Thus Railton emphasizes the undesirable epistemic habits that apriorism tends to induce: lack of openness to the possibility of new discoveries, making philosophical judgments *apriori* in the sense of *prior to, or without, substantive inquiry*, blindness to historical lessons, static outlook on concepts and knowledge, ignoring the need for change and the advantages of innovation, and so on. Hawthorne claims that the apriori-aposteriori distinction “fails to carve at the epistemological joints” and that its prospects of proving “fertile... from the point of view of serious epistemological theorizing... seem rather dim” (Hawthorne 2007: 201, 218). And Williamson argues that “although a distinction between a priori and a posteriori knowledge (or justification) can be drawn, it is a superficial one, of little theoretical significance” (Williamson 2013).²⁸

Furthermore, whereas recognition of intellect’s centrality to human knowledge has led some contemporary philosophers—e.g., Bealer (1996) and the later BonJour (1998)—to embrace apriorism, the same recognition, coupled with a clear awareness of apriorism’s limitations, has led others to opt for an intermediate position. Friedman’s “relativized a priori” (2001), for example, falls under this category. Friedman focuses on a special type of largely intellect-based principles, namely, “constitutive” principles, scientific-mathematical-philosophical principles whose role is to “provide the necessary framework within which the testing of properly

²⁷ All this will be further elaborated in Chapter 8, Section 4 and Chapter 10, Sections 5 and 9.

²⁸ (i) This citation is taken from an abstract for Williamson’s paper that does not appear in the printed version. In the printed version, Williamson says that the distinction is “of no special significance.... [It] follow[s] similarities and differences that, although genuine, are largely superficial” (Williamson 2013: 291). “[In some ways,] reliance on a distinction between the a priori and the a posteriori does more harm than good in epistemology” (Williamson: 292). “[T]he a priori—a posteriori distinction does not cut at the epistemological joints” (Williamson: 294). “We can use the qualifiers ‘a priori’ and ‘a posteriori’... if we like, but... we should not expect them to do much work in epistemology” (Williamson: 300). “Since the terms ‘a priori’ and ‘a posteriori’ are not meaningless by normal standards, some difference between a priori and a posteriori knowledge remains. But that does not rehabilitate the distinction as of great theoretical value for epistemology” (Williamson: 309).

(ii) As before, my purpose in mentioning these and other figures below is not to marshal support for my view, but to point its relation to current philosophical trends. (For a general overview of the current challenges to the apriori, see e.g., Casullo 2015.)

empirical laws is then possible" (Friedman: 83), and which "change and develop along with the continual progress of empirical and natural science, and in response to empirical findings" (Friedman: 71). Indeed, even recent supporters of the apriori, like Jenkins (2008), hold views that are in certain ways closer to those of *qualified* apriorists than of "pure" apriorists. (Jenkins's view combines an apriorist account of mathematical knowledge with an empirical grounding of mathematical concepts.)

We may conclude this section by saying that it is the view of intellect as a major yet multifaceted source of, and contributor to, veridical knowledge, operating in a variety of ways and in multiple combinations with other sources, that is characteristic of the present model.

5.4 Robust yet Non-Rigid Realism

Our model of knowledge is a realist model. Its realism is rooted in our conception of knowledge (qua knowledge) as knowledge of reality and in the importance we attribute to epistemic friction, in particular to the universal grounding-in-reality requirement. It is further based on our substantive epistemic-methodological considerations in support of NAS (e.g., the possibility that nature will surprise us by attacking the "analytic" zone of our system of knowledge), and our demonstration that a center-periphery model that grounds all branches of knowledge in reality is structurally feasible (by giving it a dynamic structure). It is also based on our considerations concerning the basic human epistemic situation, on our broad (yet non-Platonistic) conception of reality as having abstract, e.g., formal, features alongside its physical features, and on our view that humans have cognitive tools for accessing the more abstract features of reality, e.g., intellect, either alone or in cooperation with sensory perception.²⁹

Our model, however, is *realist* in a non-traditional sense. We may call its realism "*basic realism*". The attribute "basic", here, is both a *magnifier* and a *qualifier*. As a magnifier it indicates that the realism in question is substantive, robust, and universal; i.e., the model is committed to a substantive realism with respect to *all* fields of knowledge, including such fields as logic that are commonly not included in the scope of realism. This is closely connected to the model's adherence to the universal grounding-of-knowledge-in-reality requirement and its conception of the periphery as home to all disciplines (in certain contexts and times). In short, realism is basic to the principle of epistemic friction underlying the model. But "basic", as we have said earlier, is also *qualifier*. As a qualifier it indicates that the realism characterizing the model is *less rigid* and *more open-ended* than traditional forms of realism. To see what this amounts to, let us begin with what are often viewed as two fundamental commitments of realism: (A) *Existence*, and (B) *Independence*.

²⁹ This realism will serve as a basis for, and be further supported by, our account of truth and logic later in the essay.

A. *Existence*. Realism with respect to a given field of knowledge X is commonly thought to require the existence of elements in the world that X-theories are anchored in. Our model shares this commitment. By demanding that our system of knowledge as a whole, and each of its branches individually, be grounded in reality, our model, and the basic realism characterizing it, requires that all fields of knowledge satisfy a substantive existence requirement. But this existence requirement can be understood in a more or less rigid manner, and *basic* realism understands it in a relatively flexible manner. By this I do *not* mean that it compromises the grounding of our theories in reality. What I mean is that it allows flexibility in determining *what* existents are involved and *how* they are connected to linguistic expressions used in X-theories grounded in reality. To succeed in the epistemic enterprise we have to anchor each and every branch of knowledge in reality, but we are free to do so in ways that diverge from specific, preconceived templates of anchoring. To describe the difference between “rigid” and “flexible” realism more systematically, let us identify a few parameters that are relevant to this difference:³⁰

(i) *Literalness and Syntax-Ontology Parity*. A realism that demands the existence of whatever objects our theories are *literally* about—individual numbers in the case of first-order arithmetic, moral facts or properties in the case of ethics, and so on—is more rigid than a realism whose requirement that knowledge be genuinely grounded in reality does not make this demand.³¹ The literalness demand is naturally associated with another rigid demand or assumption, that of *parity of type*: the use of linguistic expressions of *syntactic* type X (e.g., singular terms) by a theory T commits T to the existence/reality of entities of *ontological* type X (e.g., individuals). Flexible realism, in contrast, requires neither literalness nor parity of type. It allows that a theory with expressions of type X satisfy the existence requirement by being anchored in the existence or reality of elements of type Y. For example, flexible realism could allow that first-order arithmetic is anchored in the reality of second-level cardinality properties, rather than in the existence of 0-level objects: numerical individuals.³² As a result, flexible realism is inclined to reject Quine’s criterion of ontological commitment, which demands sameness of type between a theory’s

³⁰ Note: in describing these parameters I will speak about realism in general, rather than about realism with respect to a field of knowledge Z. The fact that one can be a realist with respect to field Z and an anti-realist with respect to field Z’ is undisputed. But since everything I say below is compatible with this point (and can be easily relativized to Z), I will simplify the presentation by speaking about realism in general, or (following Dummett 1963) *a* realism.

³¹ In the philosophical literature, non-literalness is usually associated with anti-realism (see e.g., Gibbard 1990 and my discussion of moral truth in Chapter 8, Section 2). But in principle, realists, too, can interpret a given discourse non-literally, for example, through non-parity of types, which we have noted earlier and will return to now.

³² My typology of “orders” and “levels” in this essay is the standard one. See Chapters 8 and 10.

variables of quantification and the objects it is committed to, on the ground that it is exceedingly simplistic and unnecessarily rigid.³³

(ii) *Directness*. A realism which requires that the connection between language and reality be *direct* is more rigid than a realism that allows *indirect* connections, e.g., connections that involve intermediate objects. The latter realism would allow that an expression *e* be connected to an object *o* through an intermediate object *o**, or a number of such objects.

One form this might take involves the introduction of *posits*:

(iii) *Posits*. A realism that allows *posits* to serve as intermediaries between theory and reality is more flexible than a realism that does not. Posits (of the kind we are talking about here) are neither linguistic nor worldly; rather, they are created by us as tools for connecting theories to reality. We could, for example, anchor first-order arithmetic in second-level cardinalities (cardinality properties of properties of individuals) by introducing an intermediate level of *posits*—posited individual numbers (level 0) which are systematically connected to second-level cardinality properties in the world. This would involve a route of reference of the form: linguistic entity → posit → something in the world.

Flexibility of types (i) through (iii) is especially important for extending realism to abstract disciplines (like mathematics) in a way that avoids philosophical problems arising for more rigid forms of realism.³⁴

(iv) *Ontological Units*. A realism that limits itself to a narrow, preconceived set of ontological units (units of existence) is more rigid than a realism that does not. We are all familiar with the difficulty involved in handling mass and vague nouns (like “snow” and “bald”) with the traditional categories of individuals and sets of individuals.³⁵ A realism that allows us to incorporate new ontological units, like those offered by mereology,³⁶ is more flexible than one that does not. But the issue does not end here. There is a sense in which even mereological units are not appropriate in many fields of knowledge, from philosophy to the social sciences. Take philosophy, for example; say, Kant’s epistemology. One of the key questions of Kant’s epistemology is “What are the necessary conditions for the possibility of knowledge?”. This question is factual. It is a factual matter whether given conditions are or are not the necessary conditions for the possibility of human knowledge. As such Kant’s question is a question about reality.³⁷ But it is not clear what ontological units should be associated with conditions-for-the-possibility-of-knowledge, or whether the reality of

³³ For Quine’s criterion of ontological commitment see his (1948, 1951b, and 1953a). For discussion of objections to Quine’s criterion see e.g., Chateaubriand (2003).

³⁴ All the points made in (i) through (iii) will be further developed in Chapter 8, Section 4.

³⁵ For multiple references see e.g., Nicolas (2013) and Sorensen (2012).

³⁶ See e.g., Varzi (2003/9).

³⁷ See Subsection “Philosophy is fact-oriented” in Chapter 4, Section 1.

such conditions is fruitfully expressed in terms of units of existence at all. The same might hold for moral attributes, psychological dispositions, social and political institutions, and so on. A realism that does not enforce preconceived units of existence upon reality is more flexible than one that does.

(v) *Aboutness*. While in general, realism does and should identify the ground of a given truth with the existence (or reality) of what it is *about*, there are exceptions. Compare the truths:

(1) Snow is white

and

(2) "Snow is white" is true.

(1) is about something physical, (2) is about something linguistic. Both can be viewed as anchored in reality, yet while the anchor of (1) is what it is about, the anchor of (2) is not. Although (2) is *about* a linguistic entity, its *anchor in reality* is something physical, in fact the same physical phenomenon that anchors (1). To wit, the *linguistic* entity "Snow is white" has the *semantic* property of being true because *in the world* the *physical* stuff snow has the *physical* property of being white. A realism that recognizes exceptions to the "aboutness" principle is more flexible than one that does not.

(vi) *Bivalence*. Another way in which rigid forms of realism may differ from more flexible forms is in their attitude toward bivalence. Following Dummett (1963 and elsewhere) it is common to associate realism with bivalence, and this can add to the rigidity of a given form of realism. The *anti-realist*, Dummett says, anchors language/knowledge in reality through a standard of *verification*, which is *not bivalent*, whereas the *realist* anchors it through a standard of *truth* which, holding for all sentences (of a given language or discourse), regardless of their verifiability, is *bivalent*. Flexible realism loosens the Dummettian tie between realism and bivalence by disconnecting bivalence and truth. Our line of reasoning proceeds as follows: realism is compatible with a correspondence conception of truth. But if truth is based on correspondence, then truth significantly depends on the way the world is. And if the formal structure of the world is non-bivalent, then truth, too—hence realism—is not bivalent. We will revisit this point in Chapter 10, but briefly we can explain it in set theoretic terms as follows: a bivalent world is one in which there are exactly two possibilities with respect to a given domain of discourse D , an object o in D , and a predicate P : (i) o is a member of the extension of P in D , and (ii) o is a member of the complement of the extension of P in D . But it is possible in principle that there are more than two options with respect to D , o , and P .³⁸ In that case, truth will not be bivalent. Once this possibility is recognized, Dummett's requirement must be rejected. We acknowledge that if truth is, say, n -valent, justification is $n+1$ valent. But this has weaker consequences

³⁸ Consider, e.g., fuzzy set theory.

than Dummett's original view. In particular, bivalence as an "automatic" test of realism is renounced.³⁹

These are some of the parameters that distinguish rigid and flexible versions of realism. And the present model, with its affirmation of multiple routes from theory to reality and free movement of theories and disciplines between center and periphery, does not only favor, but actually demonstrates the structural feasibility of flexible realism.

It is important to emphasize again that by supporting *flexible* realism our model does not support a *weak* form of realism. There is no direct correlation between rigidity and strength with respect to realism. Indeed, in an important sense rigidity is an obstacle to strength. By allowing more ways of anchoring sentences and theories in reality we create resources for extending realism to fields that either have been excluded from existent versions of realism (e.g., logic) or have faced serious problems (e.g., mathematics). While rigid realism forces us to decide right away that logic, say, lies outside the realm of realism, flexible realism allows us to treat it as an open question whether, and *how*, logic is, or might be, grounded in reality. As we will see in Chapter 8, Section 4 and Chapter 10, allowing new patterns of anchoring sentences and theories in reality does indeed make it feasible to give an account of logic and mathematics that falls within the scope of realism, as well as to develop solutions to (or a strategy for avoiding) many of the problems facing traditional mathematical realism. Clearly, a realism that grounds hitherto problematic disciplines in reality in a genuine if non-traditional manner is in an important sense stronger than a realism which does not (genuinely) ground them in reality at all.

B. Independence. The second requirement associated with realism is *independence* or, as it is usually called, *mind-independence*. This independence is described in the literature in a variety of ways. Given a theory, a realist might regard its objects or its truths as (i) independent of humans' ability to discover them, (ii) independent of thought, talk, knowledge, experience, observation, awareness, perception, conception (conceivability), recognition, cognitive access by humans, (iii) independent of the existence of humans and their epistemic powers, (iv) independent of linguistic practices, conceptual schemes, theories, beliefs, etc., (v) independent of biases, preferences, social norms, ideological commitments, prejudices, personal loyalties, ambitions (of people who conduct inquiries about them), and so on.⁴⁰ But the import of these diverse descriptions is the same.

³⁹ A clarificatory note: in this essay I talk as if all branches of knowledge are bivalent. But this is intended just for the sake of simplification. All or most of the general points I make are compatible with the possibility that reality, hence truth, is non-bivalent, and they can be straightforwardly adjusted to cases in which this possibility is realized.

⁴⁰ This is a pastiche of characterizations given by Dummett (1963: 146), Boyd (1989: 6, 2002: 1), Rosen (1994: 283—(v) is taken from Rosen almost verbatim), Khlentzos (2001/11: 1, 3–4), Grayling (2001: 2–3), Jenkins (2005: 199), and others.

To appreciate the sense in which independence lies at the epicenter of realism, consider the following descriptions of this doctrine:

Realism I characterize as the belief that statements of the disputed class possess an objective truth-value, **independently** of our means of knowing it: they are true or false in virtue of a reality existing **independently** of us (Dummett 1963: 146).

Realism is a declaration of **independence**. . . . A realist view of a sentence will first cite a semantical property and then claim that the sentence has that property **independently** of *something*. Frege . . . assert[ed] that the truth value of an arithmetic statement **does not depend** on our linguistic conventions, or on the mental images that may accompany our discourse, or on our epistemic abilities to find out what the truth value of the sentence is. The “somethings” which Frege singled out in his **independence claim** are entirely typical. Realists from Plato onward, have argued that man is *not* the measure of all things. Our mental contents and abilities form a very small part of the universe. Even if our minds were much different—even if there were no sentient beings at all— $2+3$ would still equal 5 (Sober 1982: 369).

[W]hat the realist mainly claims is the right to say things like this:

Our discourse about *X* concerns a domain of fact that is *out there*. These facts obtain *anyway*, regardless of what we may think. When all goes well, inquiry in the disputed area *discovers* what is *already there*, rather than *constituting* or *constructing* its objects. Successful thought amounts to the *detection* of something *real*, as opposed to a *projection* onto the real of our own peculiar or subjective perspective. . . .

. . . For the realist . . . the target discourse describes a domain of genuine, **objective fact** (Rosen 1994: 278).

Different authors approach realism from different perspectives: some from a meta-physical perspective, others from a semantic or an epistemic perspective. The present approach is epistemic, but its view of knowledge is sufficiently broad to encompass many elements of the metaphysical and semantic perspectives as well.

It is important to distinguish between the sense in which our realism is “epistemic” and another, more prevalent, sense of “epistemic realism”. Our realism is epistemic in the sense that we are interested in the question of realism from the point of view of knowledge—as part of an attempt to develop a theory (model) of knowledge, to understand the nature of, and the constraints on, knowledge. The question is, what standard of correctness does a field need to adopt in order to provide genuine knowledge, and our epistemic realism says that it needs to adopt a standard of *truth*, where “truth” is understood in some strong sense of “correspondence”. The other sense of “epistemic realism” is that in which realism is associated with justification conditions rather than with truth conditions (in the above sense of truth). An epistemic realist in that sense might say that a theory is correct if it agrees with what our system of knowledge will say at the ideal limit of inquiry (at the ideal completion of the process of justification). This is a weaker standard than the standard of truth as (strong) correspondence, and therefore, our own epistemic

realism is stronger than epistemic realism in the second, more familiar sense.⁴¹ Accordingly, our existence and independence requirement are also stronger than that of epistemic realists of the second, more familiar kind.

However, a strong independence requirement (just like a strong existence requirement) is not the same thing as a rigid, inflexible, independence requirement. The independence requirement can be formulated in a more or less rigid manner, and accordingly be identified with a more or less rigid version of strong or robust realism. The rigid realist would say that aside from the way in which the reality of *objects created by us*, like tables and chairs, is partly dependent on us, *there is no other way* in which reality, as the measure of theories, is dependent on us.⁴² The flexible realist would recognize the role of the human mind in shaping our conception of reality, and thus say that the standard of adequacy of human theories must appeal to aspects of reality that are *largely or significantly* independent of our mode of cognizing them, but need not (and possibly cannot) appeal to aspects of reality that are *completely or absolutely* independent of our mode of cognition.

Basic realism, the realism countenanced by our model of knowledge, is a robust, yet flexible, realism. One of the advantages of such a realism is that it distinguishes between various degrees of independence. While rigid realism does not distinguish between (i) theories grounded in something that is *significantly* (though not completely) independent of us and (ii) theories grounded in something that is *completely* (or largely) dependent on us, flexible realism does. Another way to put it is that rigid realism is conducive to a false dilemma, namely the view that either knowledge is *fully (exclusively) dependent on the world* or it is *primarily dependent on the mind*. Either our theories depend *only* on the world, or we are *idealists*. The flexible realist rejects this dilemma. *Genuine* knowledge, qua *knowledge*, depends both on the world and on the mind, both on the knower and on her target, reality. The true question of realism is not whether knowledge is dependent *only* on the world (and *not at all* on the mind) but whether its dependence on the world is *sufficiently significant* (strong, deep, thorough, systematic, substantial) to justify its being *genuine* knowledge. And to say that reality, as an arbiter of our theories, needs to be significantly (but not completely) independent of the mind is to say just that. Realism, from the point of

⁴¹ One philosopher who is often mentioned in connection with epistemic realism in the second sense is Putnam during his "internal realism" phase (see, e.g., Baghramian 2008). It should be noted, however, that Putnam's views are too nuanced and subtle to be accurately described by the simple characterization given above. Throughout his career Putnam has searched for a view of realism that is right for human knowledge, and this has led him to a succession of views that attempt to overcome a variety of difficulties that arise for other views, including his own at earlier stages of his career. The views considered by Putnam, however, do not include the *basic realism* proposed here, as will be pointed out below. (In Chapter 8 we will point to similar differences between the correspondence conception of *truth* developed in this essay and the correspondence conception considered in Putnam's writings.)

⁴² This is how Miller characterizes realism in general: "The realist wishes to claim that apart from the mundane sort of empirical dependence of [man-made] objects and their properties familiar to us from everyday life, there is no *further* sense in which everyday objects and their properties can be said to be dependent on anyone's linguistic practices, conceptual schemes, or whatever" (Miller 2002/14: 2).

view of the flexible realist, concerns the extent to which knowledge is grounded in reality, not the extent to which nothing else takes any part in grounding knowledge. Indeed, the exclusiveness requirement, according to the flexible realist, is not just unnecessary, but *distorts* realism. To require the absence of any dependence of knowledge on the mind is to close our eyes to a critical constituent, a *sine qua non*, of human knowledge.

As we noted earlier, however, flexible realism is not *ipso facto* a weak realism. On the contrary. Rigid realism is likely to exempt large parts of our system of knowledge—namely all those that cannot satisfy its rigid independence requirement—from *any* independence requirements (just as it is likely to exempt them from any existence requirements), and in this sense it is weaker, from an epistemic perspective, than a flexible realism that imposes a *significant* independence requirement on *all* fields of knowledge. Basic realism is strong in this way. It sets a *strong* independence requirement on our system of knowledge, demanding that *all* branches of knowledge be *substantially* grounded in a reality *significantly independent* of the mind. But, in requiring highly diverse disciplines to be grounded in reality in this way, basic realism is open to the possibility of a multiplicity of ways of (genuinely) grounding knowledge in reality. Furthermore, while it requires that human knowledge depict the world correctly, it recognizes that a correct depiction of the world is partly a matter of what cognitive resources are used. Indeed, it is this combination of firmness and flexibility that makes “basic realism” *basic* in our sense. By having a more nuanced understanding of independence, a model of knowledge embracing *basic realism* is capable of offering a truer, richer, and more fruitfully demanding account of knowledge than less flexible models.

Although the basic realism characterizing our model is a strong type of realism (in the respects indicated above), it does not make any of the commitments that led Putnam (1981, 1983a, 1994c, 1999) and others to reject strong realism. The strong realism that Putnam (*rightly*) *rejects* is a particular type of strong realism—“metaphysical realism”, in his terminology. Such a realism involves commitment to a *thing in itself*, a *God’s eye view*, *nature’s having its own language*, there being *only one way* to correctly theorize about the world, and so on. None of these commitments is involved in basic realism or in any other aspect of our model.⁴³

Another distinctive characteristic of our model is its requirement that our system of knowledge, and especially philosophy, substantively investigate, and offer a substantive account of, the *interface* between mind and reality. Some philosophers—e.g., Rorty (1979) and McDowell (1994)—believe such an investigation is impossible.

⁴³ For additional clarifications see discussion of truth in Chapter 8. I should note that while our realism is not as excessive as the one Putnam calls “metaphysical realism”, it is significantly stronger than Putnam’s internal realism, at least when the latter is understood as rejecting a robust correspondence notion of truth. An evaluation of Putnam’s internal realism, however, requires further discussion of epistemic freedom, which we leave for a sequel.

I disagree. Although progress on this problem requires a greater focus on epistemic freedom than we offer in this essay, I hope that the model of knowledge, epistemic methodology, investigations of truth, and conception of logic's relation to reality developed here will make some contribution to a more thorough study of this problem in the future.

Our approach to reality, intellect, and realism deepens our differences from Quine. We have discussed some of the differences between his model and ours in Chapter 3. In the next chapter we will discuss further differences between the two models. We will then show how influential criticisms of Quine's model do not affect ours. As before, we will limit ourselves to issues that are important for understanding the theory developed in this essay.

6

Differences with Quine

6.1 Contrasts with Quine's Model

We have seen how the present model differs from Quine's in its broader and more dynamic conception of periphery and center, in its greater emphasis on the epistemic dimension of the no analytic-synthetic bifurcation thesis (NAS), in distancing itself from Quine's linguistic arguments for this thesis, in rejecting those parts of Quine's theory that suggest a one-unit holism, and more. It is also clear that the present model differs from Quine's in its open-ended conception of reality, in its view of intellect as playing a veridical, and not just pragmatic, role in knowledge, and in its approach to realism. In discussing our differences in this chapter I will focus on two themes: (1) Quine's radical empiricism as underwriting his neglect of intellect, and (2) Quine's conflicted approach to abstract knowledge: logic, mathematics, and theoretical science.

I. Quine's radical empiricism as underwriting his neglect of reason (intellect)

Quine is an avowed empiricist. He is an "enlightened" empiricist (Gibson 1988) as far as his rejection of the analytic-synthetic dichotomy and his holism are concerned, but he is a radical empiricist all the same. This, I believe, is the main source of his narrow conception of the periphery and, with it, reality. In particular, it is the source of his myopic approach to reason or intellect. Quine never considers, let alone investigates, the possibility that in addition to playing a pragmatic role in knowledge (being the seat of pragmatic considerations, so to speak), reason plays a veridical role in knowledge as well, i.e., a role in discovery, evidence, and theoretical justification. His holism allows him to assign a *residue* of veridicality to reason through its distant connections with experience, but his empiricism forbids him to contemplate the possibility that reason and experience are two equal sources of veridicality, distinct yet interconnected. Another way to put it is that when it comes to reason Quine's empiricism impedes him from following Wittgenstein's injunction "Look and see!": *Look and see* whether theoretical knowledge can be obtained without active participation of intellect in discovery and veridical justification! Instead, Quine's empiricism saddles him with a preconceived view of intellect and the resources required/available to anchor human knowledge in reality.

Quine sometimes says that his empiricism is motivated by *pragmatic* considerations, namely, by the *efficacy* of experiential evidence as compared with other types of evidence. In other places he says that science has taught us that the only cognitive route to reality is the sensory route. It is important to note, however, that Quine's "comparison group" for experiential evidence and routes to reality *does not include reason*. Instead, it consists exclusively of *irrational* routes and sources of evidence: "Homer's gods" (Quine 1951a: 44), "clairvoyance" (Quine 1973: 2, 1981c: 181), and their like. In the entire Quinean corpus there is no consideration of the one serious candidate for a non-experiential source of knowledge—*reason* or *intellect*. Not only does Quine ignore the possibility that reason is a source of knowledge all by itself, but he ignores the possibility, *suggested by his own NAS and relational holism*, that reason plays a substantial veridical role in knowledge *jointly* with experience. His yardstick for evaluating alternative sources of knowledge is limited to success in *sensory prediction* on the one hand and *pragmatic arrangements* on the other, and this blinds him to the possibility that reason makes a substantial contribution to veridicality.

It is significant to point out, in this connection, that Quine says very little about reason in his entire corpus. There is no entry for "reason" (or "intellect") in the index of any of Quine's books or collections of essays, with one exception: his reference to Leibniz's distinction between "truths of **reason**"¹ and truths of fact" in the section "Background for Analyticity" of "Two Dogmas of Empiricism" (Quine 1951a: 20). Nor is there an entry for "reason" in scholarly books about Quine (Gibson 1982, 1988, Romanos 1983, Dilman 1984, Gochet 1986, Hookway 1988, Hylton 2007) or in collections of essays about Quine (Shahan and Swoyer 1979, Hahn and Schilpp 1986, Barrett and Gibson 1990).² It should be noted that in one place Quine does allow a somewhat greater role to reason in knowledge:

The contribution of **reason** cannot be viewed as limited merely to conceptualizing a presented pageant of experience and positing objects behind it; for this activity reacts, by selection and emphasis, on the qualitative make-up of the pageant itself in its succeeding portions (Quine 1952: 224).

But even here reason's role is quite limited, and in any case, what he says here seems to have had no impact on his overall approach to reason.

I should emphasize that my issue with Quine is not the role of empirical evidence in science—any viable model of knowledge must acknowledge the central role of such evidence in scientific knowledge. Nor is it the positive role of sensory experience in discovery. The issue is the uncritical acceptance of experience as an *exclusive* source of discovery and evidence. We may say that while Quine is an iconoclast with respect to foundationalism and the analytic-synthetic distinction, he is a conformist and a

¹ For **boldface** within citations see fn. 3 in Chapter 1.

² "Reason" is mentioned in an article on Quine by Putnam (1983c), but this article is devoted to criticisms of Quine's neglect of reason.

traditionalist with respect to the basic empiricist principles of evidence and information. Like most traditional empiricists he believes that there is exactly one open channel between our system of knowledge and reality, the sensory channel, and like the logical positivists he limits the other constituents of knowledge to the pragmatic and conventional.

Now, someone might point out that it is common for scientists to claim that experience is the final arbiter in science, that theory must conform to observation. This might be a significant part of Quine's position, for it seems to create an asymmetry between reason and experience.

This is a fair point, but my criticism of Quine is directed at his *philosophical theory* of scientific knowledge, an endeavor that requires more than just repeating what scientists say. The fact that practicing scientists emphasize the crucial role of experiment in their work does not rule out the possibility that reason plays a central role in it as well.

Quine's radical empiricism is expressed time and again in his writings, in such statements as:

[I]n our knowledge of the external world we have **nothing** to go on **but surface irritation** (Quine 1954b: 230).

Quine is an empiricist through and through:

My basic position **early and late** is **empiricism**, and hence prediction as touchstone (cited in Dreben 1992: 308).

His theory is largely motivated by problems internal to empiricism, as can be seen in his two flagship works, "Two Dogmas of Empiricism" and "Epistemology Naturalized". Rather than "Two Dogmas of *Knowledge*", the first paper is titled "Two Dogmas of **Empiricism**", and the problems it deals with are presented as problems *within* empiricism:

Modern empiricism has been conditioned in large part by two dogmas. One is a belief in some fundamental cleavage between truths which are *analytic*, or grounded in meaning independently of matters of fact, and truths which are *synthetic*, or grounded in fact. The other dogma is *reductionism*: the belief that each meaningful statement is equivalent to some logical construct upon terms which refer to immediate experience. Both dogmas, I shall argue, are ill-founded (Quine 1951a: 20).

Quine renounces both dogmas, but this renouncement is, for him, a way of *preserving* empiricism rather than *rejecting* it. In delineating his alternative model, Quine clearly depicts himself as an empiricist:

As an empiricist I continue to think of the conceptual scheme of science as a tool, ultimately, for predicting future experience in the light of past experience (Quine: 44).

The theoretical goal of science, according to Quine, is to find an "efficacious . . . device for working a manageable structure into the **flux of experience**" (Quine: 44), and one of the main goals of his model is to show that this goal is best attained by rejecting the

two staples of earlier empiricist models mentioned above: the analytic-synthetic thesis and the reductionist thesis. Quine concludes his paper in an empiricist vein as well: “Each man is given a scientific heritage plus a continuing barrage of **sensory stimulation**”, and his goal is to “warp . . . his scientific heritage to **fit his continuing sensory promptings**” (Quine: 46).

The program advanced by Quine in his second leading epistemic paper, “Epistemology Naturalized”, is also empiricist. Thus, speaking about this program in his (1990) book, he says:

The most notable norm of **naturalized epistemology** actually coincides with that of traditional epistemology. It is simply the watchword of **empiricism**: *nihil in mente quod non prius in sensu*. . . . [O]ur information about the world comes **only** through **impacts on our sensory receptors** (Quine 1990: 19).

But Quine’s choice of empiricism over other epistemic outlooks, as he presents it, is uncritical. Although occasionally he suggests a pragmatic motivation for empiricism or mentions (in passing) that science itself shows that our only access to the world is sensory, for the most part he simply takes empiricism as given. Significantly, not once in his entire corpus does Quine either question empiricism or offer a theoretical justification for it.

Dreben feels that many philosophers have failed to see the crucial role played by empiricism in Quine’s philosophy:

The essential point that is missed again and again about Quine is that his basic position has always been what he calls “**empiricism**”; **prediction is its touchstone** (Dreben 1992: 308).

But some have recognized the centrality of empiricism to Quine’s theory (e.g., Davidson 1981, Gibson 1982 and 1988, Hookway 1988, and Hylton 2007). And a few emphasize its problematic nature (e.g., Davidson and Hookway). From my own perspective, the main problem with Quine’s empiricism is its *philosophical narrowness* and *one-sidedness*.

The model delineated in the present essay acknowledges the crucial role of empirical methods in science. But it is a leap to get from this acknowledgment to radical empiricism, and this leap we do not make. We are interested in understanding the nature of human knowledge in all its complexity, including the broad array of roles that intellect plays in it. This difference has sweeping ramifications for many aspects of the two models.

Another source of differences between the two models, though of a secondary importance, is the “linguistic turn”.

Quine grew up, philosophically, in the first half of the twentieth century, and he remained, at least to some degree, a product of the “linguistic turn” which dominated much of English-speaking philosophy during that period. This is reflected both in his interests and in his argumentation. We have already seen that, in spite of the fact that Quine’s rejection of the analytic-synthetic dichotomy led him to a new *epistemic*

model, his arguments against it in “Two Dogmas of Empiricism” were, for the most part, linguistic.

The new model, as we have pointed out above, regards Quine’s linguistic arguments as irrelevant. It supports NAS based on distinctly epistemic considerations, like the “Maginot line” of knowledge, the universality of the veridical-justification requirement, and the dynamic nature of our system of knowledge and its concepts (see Chapter 3). More generally, it is based on the principles of epistemic friction and freedom (as formulated in Chapter 1, Sections 2–3) and those of foundational holism (delineated in Chapter 2, Section 2). All these support an epistemic “non-bifurcation” thesis (NB), a thesis that says (roughly and informally):

- (NB) Our system of knowledge is not bifurcated into units substantially grounded in fact (reality) and units exclusively grounded in something other than fact. All units of knowledge, qua units of knowledge, are substantially grounded in both fact and mind.

And NB, in turn, leads to NAS, which can roughly be formulated as follows:

- (NAS) Our sentences are not divided into sentences whose truth depends on fact (reality) and sentences whose truth depends only on something other than fact (mind, language, concepts, conventions). The truth of all sentences depends both on fact and on the mind.

In his (1991) retrospect of “Two Dogmas”, Quine describes his rejection of NAS as motivated by his antipathy to “mentalistic notions”, “distrust of mentalistic semantics”, “[un]abetted extensionalism”, rejection of “propositional functions”, and strict “aloofness from intensions and mentalism” (Quine: 265–7). The new model shares *none* of these attitudes either.

II. Quine’s conflicted approach to abstract knowledge: logic, mathematics, and theoretical science

Abstract knowledge in Quine’s model is largely instrumental (as many commentators—e.g., Gibson 1982: 167 and Esfeld 2001: 204—have noted); in our model it is informative in its own right. Highly abstract units of knowledge in Quine’s model can be grounded in reality only nominally, by partaking in the general grounding of our system of knowledge as a whole; those same units can, in our model, be grounded in reality also as independent branches of knowledge whose task is to account for features of reality that are largely inaccessible to sensory perception.

Quine’s instrumentalist approach to abstract knowledge is, however, countered by his realism and NAS.³ Both Quinean realism and NAS place abstract areas of

³ In discussing Quine’s instrumentalism, Hylton (2007: 18–23) notes that Quine’s instrumentalism is also mitigated by his view that *all* objects are *posits*. Unlike other instrumentalists, Quine does not contrast abstract and non-abstract objects.

knowledge on a par with empirical areas, and this means that the two equally require a grounding in reality. But Quine's empiricism denies him the requisite resources for such a grounding. This leads to tensions in Quine's approach to specific abstract disciplines, analogous to those we have observed between his NAS and CP. None of these tensions arise in our model.

Logic. NAS's affirmation of the substantial factuality of all fields of knowledge commits its adherents to the view that as a genuine field of knowledge, logic, too, is substantially factual. And indeed Quine boldly declares:

Logical theory, despite its heavy dependence on talk of language, is... **world-oriented** rather than language-oriented (Quine 1970/86: 97. Cited earlier).

More tentatively, he suggests that:

[L]ogical truth... may... depend on... features of the world... that our language reflects in its grammatical constructions rather than its lexicon (Quine: 95).

But Quine's narrow empiricism prohibits him from viewing logic as world-oriented in a *serious* sense. If "the world can be evidenced only through stimulation of our senses" (Quine 1955: 252), how can logic, whose claims are too abstract and too general to be justified based on such evidence, be substantially world-oriented or grounded in reality? Logic can simplify other theories which are seriously world-oriented, but it cannot be seriously world-oriented itself. It is not surprising, therefore, that Quine makes light of his claim that logic is world-oriented:

[O]ne could... maintain that the... matter of Tom's being either mortal or not mortal is due... to pervasive traits of nature... **if one could make sense of the issue at all** (Quine 1970/86: 15).

And again:

Is logic a compendium of the broadest traits of reality, or is it just an effect of linguistic convention?... [This question] **has proved unsound; or all sound, signifying nothing** (Quine: 96)

One might try to resolve the conflict by suggesting that while Quine cannot regard logic as *epistemically* factual he can, and does, regard it as *metaphysically* factual. (This view would appeal to those who draw a sharp distinction between Quine's epistemology and metaphysics.) But this solution will not work, since it is the metaphysical version of the claim that logic is factual that Quine ridicules in the above citations.

Furthermore, although Quine declares that logic is world-oriented rather than language-oriented, this claim is necessarily deflated by his approach to truth. The worldly orientation of logic, according to Quine, is due to its connection with truth: it is "the **truth predicate**" that "makes" logic "world-oriented rather than language-oriented" (Quine: 97). But to the extent that Quine's conception of truth is deflationist, logic's connection to the world is also deflationist.⁴ This deflationist approach

⁴ Deflationism will be discussed in Chapter 7.

is reflected in his claims that logic is “vacuous” and “obvious”: “Logic is true... **vacuously**, it is true by virtue of anything and everything” (Quine: 97). The “canon—‘**Save the obvious**’—is sufficient” to justify logic (Quine: 82). My own view of logic could not be more different.⁵

The inner tension in Quine’s approach to logic is on display even in passages which emphasize the similarity between logic and science. On the one hand, as a supporter of NAS Quine is led to say that a revolution in logic is possible and, moreover, would be similar in kind to a revolution in science:

Revision even of the logical law of the excluded middle has been proposed... [by some practitioners of] quantum mechanics; and what difference is there in principle between such a shift and the shift whereby Kepler superseded Ptolemy, or Einstein Newton, or Darwin Aristotle? (Quine 1951a: 43).

On the other hand, it is hard for Quine, as an empiricist, to say that a revolution in logic would be *as factual* as a revolution in science. His way out is to turn the tables and say that a revolution in science is *as pragmatic* as a revolution in logic. Thus the full citation is:

Revision even of the logical law of the excluded middle has been proposed **as a means of simplifying** quantum mechanics; and what difference is there in principle between such a shift and the shift whereby Kepler superseded Ptolemy, or Einstein Newton, or Darwin Aristotle? (Quine: 43).

That is to say: Quine equates a revolution in logic with a revolution in empirical science not by affirming the *factuality* of logic but by affirming the *pragmaticality* of empirical science.

Another manifestation of this tension is found in Quine’s justification of his own theoretical preferences in logic. In spite of declaring that logic is factual—something that calls for a veridical justification of his choice of logic (or at least a justification with a significant veridical component)—Quine’s justification of his logical preferences, namely, of preferring standard first-order logic over all other logics, is *pragmatic*:

The **rewards** of staying within the bounds of standard grammar [i.e., the grammar of standard first-order logic] are great. There is **extensionality**.... There is, more generally speaking, the **efficiency** and **elegance** of the logic of truth functions and quantification. There is its **completeness**.... There is the impressive **concurrence** of [a number of] definitions of logical truth (Quine 1970/86: 79).⁶

⁵ For a critical discussion of the view that logic is *obvious* and a substantivist treatment of logic see Chapter 9, Section 2 and Chapter 10.

⁶ For Quine, in (1970/86), the advantages of extensionality, completeness, and the concurrence of multiple definitions of the same notion are *pragmatic*.

Furthermore, a recurrent theme in Quine's rejection of deviations from standard logic is the pragmatic principle of "minimum mutilation":

[T]hree-valued logic... runs counter to a generally sound strategy which I call the **maxim of minimum mutilation** (Quine: 85).

Whatever the technical merits of the case [for a logic of quantum mechanics], I would cite again the **maxim of minimum mutilation** as a deterring consideration (Quine: 86).

If revisions are seldom proposed that cut so deep as to touch logic, there is a clear enough reason for that: the **maxim of minimum mutilation** (Quine: 100).

Other considerations in favor of standard logic are also pragmatic or at least non-factual:

[L]et us not underestimate the price of a deviant logic. There is a serious loss of **simplicity**. . . . And there is a loss, still more serious, on the score of **familiarity** (Quine: 86).

Intuitionist logic lacks the **familiarity**, the **convenience**, the **simplicity**, and the **beauty** of our logic (Quine: 87).

It is quite remarkable that not even once does Quine appeal to considerations of *correctness* (*incorrectness*) in justifying his choice of logic or in rejecting any deviation from his chosen logic. Given his empiricist notion of evidence, this is understandable. Although NAS says that all units of knowledge are substantially factual, empiricism forces him to treat logic as no more than nominally factual, i.e., as primarily pragmatic and conventional.

An important consequence of Quine's pragmatist approach to logic is his inability to pursue a *veritable foundation* for logic. Given his claim that logic is "world-oriented", such a foundation would have to provide a substantive account of the way the logical laws are anchored in reality. And like any other foundational account, it would also have to provide (i) a theoretical justification of the current, or new, criteria, or definitions, of logical truth and consequence, (ii) a general, precise, and informative demarcation of logical terms or operators (or else an explanation of why such a demarcation is not needed), (iii) an explanation and justification of the scope of logic, and so on. But Quine's essay on the nature of logic (1970/86) does not—and, given Quine's conflicted attitude toward logic, cannot—provide any of these things.

The present model, in contrast, is free to do all these things. It does not suffer from any of the tensions characterizing Quine's attitude toward logic, it is fully committed to the grounding of logic in both world and mind, it has no extraneous ideology that would prevent it from pursuing a foundation for logic, and it has a foundational-holistic methodology that makes such a project feasible.⁷

⁷ See Chapters 9 and 10.

Mathematics. Quine is also pulled in opposite directions with respect to mathematics. While his criterion of ontological commitment and NAS pull him toward *mathematical realism*, his empiricism pulls him toward *mathematical instrumentalism*. Thus on the one hand Quine says:

[M]y ontology admits... the full complement of extensional universals, mathematical and otherwise (Quine 1986a: 315).

I...espouse... a realism of universals... Mathematics... is up to its neck in universals (Quine 1981d: 182).

I am a Predicate and Class realist, now as of yore; a deep-dyed realist of abstract universals (Quine: 184).

But on the other hand he advocates an instrumentalist approach to mathematical ontology:

As an empiricist I... think of... the abstract entities which are the substance of mathematics [as]... expedit[ing] our dealings with sense experiences (Quine 1951a: 44–5).

The reason for admitting numbers as objects is precisely their efficacy in organizing and expediting the sciences. The reason for admitting classes is much the same (Quine 1960: 237).

Limited to physical objects though our interests be, an appeal to classes can... be instrumental in pursuing those interests. I look upon mathematics in general in the same way, in its relation to natural science (Quine 1981b: 15).

My view of pure mathematics is oriented strictly to application in empirical science... [W]e admit the extras to simplify our computations and generalizations (Quine 1986b: 400).

Abstract objects have... proved indispensable to natural science - thus numbers, functions, classes (Quine 1990: 237).

In other words, mathematics is valuable not in providing bona fide mathematical knowledge, but in making physical, biological, psychological, economic, and other types of empirical knowledge possible. This is Quine's *indispensability argument* for mathematics. Putnam's formulation of this principle rightly emphasizes its *instrumentalist* or *pragmatist* nature:

[I]n the last analysis it is the utility of statements about mathematical entities for the prediction of sensory stimuli that justifies belief in their existence (Putnam 1994b: 245).

Quine's justification of not dispensing with "higher" mathematics is also pragmatic:

As for inapplicable parts of mathematics, say higher set theory, ... [w]e... keep their sentences as meaningful... because they are built of the same lexicon and grammatical constructions that are needed in applicable mathematics. It would be an intolerably pedantic *tour de force* to gerrymander our grammar in such a way as to account the inapplicable flights ungrammatical while preserving the applicable part (Quine 1991: 269).

The approach to mathematical knowledge advocated here is quite different. While mathematics has close connections with science, its theories are informative in their

own right. The underlying idea, in its simplest form, is expressed by Gödel: “That the flower has *five petals* is as much a part of **objective reality** as that its color is *red*” (Wang 1987: 202).⁸ Objects and structures of objects in the world have formal features along with physical, biological, and other kinds of features, and a primary task of mathematics is to study these features, discover the laws governing them, and systematize the results by informative theories.⁹

This does not mean that mathematics plays no instrumental role in knowledge. On the contrary, like all other disciplines, mathematics plays both a direct informative role and an auxiliary instrumental role. (Mathematical theories may alternate these roles, as well as perform both roles during a single period.) But even in their non-directly factual capacity, mathematical theories have a factual element. Indeed, even when performing an “if-thenist” role,¹⁰ mathematical theories make factual claims. They make the claim that if *the world* (or one of its structures) *is* as the theory’s *axioms* say it *is*, then *it is* also as the theory’s *theorems* say it *is*. Such a claim requires a veridical, and not just a pragmatic, justification.¹¹

Abstract Scientific Knowledge. Similar tensions to those seen in Quine’s approach to logic and mathematics are also seen in his approach to theoretical science (although perhaps to a somewhat lesser degree, due to the greater proximity of science to sensory experience). Thus, sometimes Quine describes the goal of science in realistic terms:

The general task which science sets itself is that of specifying how reality “really” is: **the task of delineating the structure of reality** (Quine 1954b: 232).

But his empiricism leads him to an instrumentalist conception of the scientific goal:

As an empiricist I...think of the conceptual scheme of science as a **tool**, ultimately, for **predicting** future experience in the light of past experience (Quine 1951a: 44. Cited above).

[T]he purpose of concepts and of language is **efficacy** in communication and in **prediction**. Such is the ultimate duty of language, **science**, and philosophy (Quine 1950: 79).

Our talk of external things...is just a conceptual apparatus that **helps us to foresee** and **control** the triggering of our sensory receptors in the light of previous triggering of our sensory receptors (Quine 1981b: 1).

Claims like these led McGinn to say:

I cannot see how Quine’s realism is ultimately to be squared with his relativistic instrumentalism (McGinn 1983: 241).

⁸ Wang says that Gödel was “very fond of [this] observation” which “he attribute[d] to Bernays” (Wang 1987: 202).

⁹ For a development of this approach see the account of mathematics in Chapter 8, Section 4.

¹⁰ See Russell (1903/38, “Introduction to the Second Edition” and Chapter 1) and Putnam (1967a, Section 3).

¹¹ For a similar point see discussion of the factual element in conventionalism in Chapter 9, Section 2.

In the end, empiricism seems to have the upper hand in Quine's conception of science. Even when he describes the aim of science as "understanding" (Quine 1990: 2, 20), Quine's understanding of "understanding" is narrowly empiricist. Quine regards observation statements and observational conditionals¹² as exhausting the worthwhile content of a scientific theory. His ideal scientific theory is *not* one that identifies and explains the *basic forces, mechanisms, and nomic regularities, or laws, of nature*, but one that provides an *effective enumeration of all observations*. Ideal science consists of an axiomatized collection of all observation conditionals which, given appropriate conditions, logically implies all observations. The best theory, therefore, is one whose *abstract content* is "**dispensed with**" (Quine 1975: 323) or, better still, a theory with no abstract content to dispense with at all: a theory that "tightly encompass[es]" observation conditionals "**without theoretical foreign matter**" (Quine: 324). And although Quine recognizes that in practice this ideal is unachievable, he believes it accurately captures the whole information conveyed by a given theory: "The empirical content of a theory formulation"—which for him is the whole information it conveys—"is summed up in the observation conditionals that the formulation implies" (Quine: 323). In contrast, the abstract portion of a scientific theory is "**trumped-up matter, or stuffing**, whose **only** service is to **round out the formulation**" of observation statements (Quine: 324); it is no more than "a **device for remote control and for mass coverage**" of "**observation conditionals**" (Quine: 324).

Quine's radically empiricist view of scientific knowledge is also expressed by his well-known bias toward observation statements: observation statements, unlike abstract theoretical statements, have "an empirical content all [their] own and wear... it on [their] sleeve[s]" (Quine 1969a: 89). They are "separately susceptible to tests of observation", in spite of "shar[ing] much of the vocabulary of the more remotely theoretical statements" (Quine 1975: 314).

It has often been commented that this is an exceedingly narrow view of science. Thus Putnam says:

[L]eaving out of science all mention of what cannot be seen with the naked eye would be leaving out just about *all* of science. (Putnam 1965: 233).

The new model regards the primary goal of science as *knowledge* and *understanding*, and it views *nature*, i.e., reality, or the physical aspect of reality (rather than observations), as the target of that knowledge/understanding. It treats abstract, theoretical content as factual, and it deems observation statements incapable of fulfilling the goal of a correct, informative, and explanatory account of nature all by themselves, even under ideal conditions.

Considering Quine's model as a whole, we can say that its traditional empiricism (or the traditional elements of its empiricism) neutralize, and in certain respects even

¹² Universally quantified conditionals whose consequent and antecedent are instantiated, respectively, by observation statements and their boundary conditions (Quine 1975: 317–18).

reverse, the considerable progress made by its more innovative principles: holism and NAS. While NAS and relational holism offer a new, non-dogmatic approach to knowledge in general and abstract knowledge in particular, Quine's empiricism takes us back to a more traditional, dogmatic approach.¹³ Like many empiricists Quine believes that Humanity is capable of only one type of contact with reality, namely, sensory contact, and this leads him to allow no more than a remote link between abstract knowledge and reality, insufficient for a serious grounding of such knowledge in its own right.

Indeed, this empiricism threatens to undermine Quine's goal of overcoming the traditional limitations of Carnap's model (as he, Quine, understands it). The central axis of Quine's model, namely:

Periphery \Rightarrow *Intermediate Zone* \Rightarrow *Center*,

partially, yet significantly, mirrors the central axis of Carnap's model:

Outermost Synthetic Zone \Rightarrow *Innerness Synthetic Zone* \Rightarrow *Analytic Zone* \cup *Metascience*.

The correlation is both extensional and normative:

<i>Extensional Correlation:</i>	<i>Quine's Model</i>	<i>Carnap's Model</i>
<u><i>Math, Logic, Language:</i></u>	<i>Center</i>	<i>Metascience / Analytic Zone</i>
<u><i>Experiential Science:</i></u>	<i>Periphery</i>	<i>Outermost Synthetic Zone</i>
<u><i>Theoretical Science:</i></u>	<i>Intermediate Zone</i>	<i>Innerness Synthetic Zone</i>
<i>Normative Correlation:</i>	<i>Quine's Model</i>	<i>Carnap's Model</i>
<u><i>Pragmatic-Conventional (P) Norms:</i></u>	<i>Center</i>	<i>Metascience / Analytic Zone</i>
<u><i>Veridical, Experiential (V/E) Norms:</i></u>	<i>Periphery</i>	<i>Outermost Synthetic Zone</i>
<u><i>Combination of V/E & P Norms:</i></u>	<i>Intermediate Zone</i>	<i>Innerness Synthetic Zone</i>

Figure 6.1

I should note that there are some exceptions to these correlations (for example, concerning philosophy's place in the two models), and there is also a difference in sharpness between the respective divisions. Nevertheless, epistemically, the correlations are quite striking.

Summing up our discussion of Quine's approach to abstract knowledge, what we see again and again is that Quine's model, like other empiricist models, is incapable of satisfying the universal friction requirements as they apply to this knowledge. Its narrow, exclusively sensory interface with reality rules out access to any non-sensory features of reality, including, most significantly, those features that are the subject matter, or ultimate target, of the abstract branches of knowledge. Quine's model allows abstract units to be grounded in reality only by evidence pertaining to the

¹³ We may say that Quine takes away with one hand what he gave us with the other.

claims of *other* units, units of an altogether different kind, namely, experiential units which cannot capture more than a *fraction* of the content of abstract units. As a result, only a modicum of abstract knowledge can be grounded in reality. Abstract units of knowledge are no more than tools for organizing, explaining, and/or predicting observations, and their justification is predominantly pragmatic and only minimally veridical.

There are other differences between Quine's model and ours, concerning truth, philosophy, and more. But these can be left for another discussion. Our next task is to show that the new model is immune to many of the familiar criticisms of Quine's model.

6.2 Immunity to Criticisms of Quine's Model

Quine's model has been subjected to multiple criticisms. How does the present model stand vis-a-vis these criticisms? In many cases, it is unaffected. This is because many of the criticisms target elements of Quine's model that are either explicitly rejected by, or have nothing to do with, the present model. In some cases, it is the non-Quinean elements incorporated into the new model that spare it from such criticisms.¹⁴ To see what criticisms the new model is immune to and why, it is useful to divide them into three categories:

- (a) *Criticisms of aspects of Quine's model that were explicitly rejected, withdrawn, or changed by the new model:*
 - (i) Criticisms of Quine's model due to the inner conflict between NAS and CP (e.g., Dummett 1973, 1973/81). [*New model: conflict resolved. Chapter 3*]
 - (ii) Criticisms of Quine's extreme empiricism (e.g., Davidson 1981). [*New model: extreme empiricism rejected. Chapters 3–5 and Chapter 6, Section 1*]
 - (iii) Criticisms of Quine's one-unit holism (e.g., Hofstadter 1954, Dummett 1973, 1973/81, Glymour 1980). [*New model: one-unit holism rejected. Chapter 2, Section 2*]
 - (iv) Criticisms of Quine's narrow/inadequate conception of reason's role in knowledge (e.g., Putnam 1983c). [*New model: Quine's narrow conception of reason rejected. Chapters 3, 5, and Chapter 6, Section 1*]
 - (v) Criticisms of Quine's rejection of meanings, mental entities, modality, intensionality, etc. (e.g., Marcus 1988, Hanna 2001). [*New model:*

¹⁴ Of course, with respect to each of these criticisms there is the question whether they apply to Quine's model in the first place—whether the reason they do not apply to the new model holds for Quine's model as well. And with respect to some of the criticisms (though by no means all) this might be the case. But to avoid unnecessary digressions, I will limit myself to the question of whether the new model is immune to these criticisms.

nothing in it involves a negative attitude toward any of these. Modalities play an active role in the model. Chapter 10]

- (vi) Criticisms of Quine's deflationism with respect to truth (e.g., Gupta 1993a, 1993b). [*New model: anti-deflationist. Chapter 7. See also commitment to substantiveness in Chapter 1, Section 2*]
- (vii) Criticisms of radical aspects of Quine's naturalism (e.g., Stroud 1981, 1984, Kim 1988, Bealer 1987, BonJour 1994). [*New model: rejects Quine's radical empiricism, and this includes the radical aspects of Quine's naturalism. Chapters 3 and 5*]
- (viii) Criticisms of Quine's radical anti-apriorism (e.g., Bealer 1987, BonJour 1994, Friedman 2001). [*New model: Rejects Quine's radical anti-apriorism; assigns central role to intellect in discovery and veridical justification; endorses quasi-apriority. Chapter 5, Section 3, Chapter 10, Section 5*]
- (b) *Criticisms of Quinean principles/arguments that play no role in the present model:*
 - (i) Criticisms of Quine's linguistic arguments for NAS (e.g., Grice and Strawson 1956, Hanna 2001, Russell 2008). [*New model: support of NAS is not based on, and has nothing to do with, Quine's linguistic arguments, or, indeed, any of Quine's arguments in Sections 1–4 of "Two Dogmas". New model's support of NAS is largely based on epistemic considerations. Chapter 3, Chapter 5, Section 4, and Chapter 6, Section 1*]¹⁵
 - (ii) Criticisms of Quine's NAS based on its relation to his indeterminacy thesis (e.g., Boghossian 1996)¹⁶. [*New model: support of NAS has nothing to do with Quine's indeterminacy thesis. Chapter 3, Chapter 5, Section 4, and this chapter, Section 1. As noted next, the indeterminacy thesis plays no role in the new model*]
 - (iii) Criticisms of Quine's theses of indeterminacy, ontological relativity, etc. (e.g., Chomsky 1968). [*New model: none of Quine's theses of indeterminacy, ontological relativity, or their cognates plays any role in the model.*]¹⁷
- (c) *Criticisms of Quinean principles that call for a response by the new model:*
 - (i) Some (partly linguistic) criticisms of Quine's NAS (that do not focus on Quine's own linguistic arguments for this thesis) (e.g., Grice and Strawson 1956).
 - (ii) Some criticisms of Quine's holism (e.g., Grünbaum 1960, Friedman 2001).

¹⁵ For Grice and Strawson's criticisms that do not relate to Quine's linguistic arguments, see Subsection I.

¹⁶ It should be noted that although Boghossian is associated with a new defense of the analytic-synthetic distinction, he represents a large group of philosophers whose attitude toward the distinction is complex: they accept some aspects of the distinction but reject others. Boghossian, in particular, rejects those aspects of the analytic-synthetic distinction that were the target of our own criticism ("analyticity 1" in his terminology). His criticism is focused on the relation between Quine's NAS and his indeterminacy thesis. Since neither Quine's indeterminacy thesis nor its negation play any role in the new model, Boghossian's criticism does not affect it.

¹⁷ I should note, however, that I do plan to discuss these issues in connection with epistemic freedom.

Clearly, our model is immune to all the criticisms falling under (a) and (b). Let us turn, then, to the criticisms falling under (c).¹⁸

I. Criticisms of Quine's NAS: Grice and Strawson

One well-known criticism, or cluster of criticisms, of Quine's NAS that might be thought to apply to our model as well appears in Grice and Strawson (1956). Grice and Strawson's criticism consists of three parts or claims:

A. The most influential claim is that existence of a significant analytic-synthetic distinction in philosophy cannot be denied:

[Quine] declares, or seems to declare, ... that [the analytic-synthetic distinction] is altogether illusory ...; so his rejection of it would seem to amount to a denial of its existence (Grice and Strawson: 142).

But this denial is unwarranted:

Is there such a presumption in favor of the distinction's existence? Prima facie, it must be admitted that there is. [There is the] philosophical tradition... But there is no need to appeal only to tradition; for there is also present practice. We can appeal, that is, to the fact that those who use the terms "analytic" and "synthetic" do to a very considerable extent agree in the applications they make of them. They apply the term "analytic" to more or less the same cases, withhold it from more or less the same cases, and hesitate over more or less the same cases. This agreement extends not only to cases which they have been *taught* so to characterize, but to new cases. In short, "analytic" and "synthetic" have a more or less established philosophical *use*; and this seems to suggest that it is absurd, even senseless, to say that there is no such distinction. For, in general, if a pair of contrasting expressions are habitually and generally used in application to the same cases, *where these cases do not form a closed list*, this is a sufficient condition for saying that there are *kinds* of cases to which the expressions apply; and nothing more is needed for them to mark a distinction (Grice and Strawson: 142–3).

If this is right, then any model, including ours, must recognize the existence of this philosophical distinction.

B. Grice and Strawson's second claim is that the existence of an analytic-synthetic distinction is compatible with Quine's model of knowledge. In particular, the distinction is compatible with Quine's universal revisability principle and his holism (or more specifically his "confirmation holism", which is closely related to what we

¹⁸ The criticisms mentioned earlier are quite numerous, but they do not include *all* the criticisms that have been directed at Quine's model. Nor do they include all the possible criticisms of our model. What I do in this section is show, within reasonable space limits, that our model is immune to many criticisms of Quine's model (as they now stand), and respond to a couple of criticisms that might be viewed by some as applying to our model along with Quine's.

have called “relational holism”). Concerning revisability, for example, Grice and Strawson say that all the existence of an analytic-synthetic distinction requires is that after a conceptual revision (possibly in response to experience) is made, there are still (or once again) necessary connections between some concepts, connections that give rise to analytic truths.

C. The third part of Grice and Strawson’s criticism consists of the claim that Quine’s rejection of the analytic-synthetic distinction threatens a host of related distinctions and concepts that are important for philosophy. These include the distinctions between the *necessary* and the *contingent*, the *apriori* and the *empirical*, truths of *reason* and truths of *fact*, having the *same meaning* and having *different meanings*, having the same *extension* and having the same *intension*, being a *correct translation* and being an *incorrect translation*, being *meaningful* and being *meaningless*, etc., as well as concepts like *believing* and *understanding*. *Analyticity* is also closely connected to *sentence significance*, *semantical rule*, *definition* (in a broad sense), and so on. These concepts and distinctions play important roles in philosophy, our culture, and our everyday lives, and as such should not be excised. (Grice and Strawson: 142–51)

As far as the new model is concerned, my response to Grice and Strawson’s claims is as follows:

(a) First, the claim that the analytic-synthetic distinction exists is largely descriptive while our model is largely normative, and there is no conflict between the claim that a given distinction is being drawn in practice and the claim that it is philosophically undesirable or counter-productive. So, to the extent that their claim is descriptive, it does not apply to our model.

Second, to the extent that their claim has a normative element as well, it is a limited normative element, and one that is also compatible with our model. Let me explain: our model requires a *productive balance* between the *fixity* and *non-fixity* of concepts. More specifically, it is committed to the view that conceptual change is, and should be, balanced by conceptual stability. While in principle no concept is held fixed once and for all (i.e., any concept can, in principle, become unfixed and thereby open to factual or theoretical influences), the stability of our system, and indeed our very ability to construct, examine, and revise theories, require that at any given time some concepts be held fixed. Now, since this view allows that some, and even many, concepts remain fixed for extended periods of time, the model can acknowledge that, and explain why, an analytic-synthetic distinction arose in the first place, and why it continues to “exist”.

This, however, does not undermine our objection to a *full-fledged* analytic-synthetic distinction in philosophy. Recognizing that a certain measure of conceptual stability is both natural and advantageous does not conflict with realizing that excessive stability—which is tantamount to immobility—is disadvantageous. Unchecked conceptual immobility blinds us to the reality, desirability, and indeed

inevitability of conceptual change, including *fact-* or *theory-driven* conceptual change. Our model rejects the assumption that there is a permanent and absolute bifurcation of statements into those whose *truth value* is, *in principle* and *in the long run*, *exclusively* a matter of *meaning* and those whose truth value is, in principle and in the long run, a matter of fact. It objects to inducing a radical bifurcation between factual and non-factual knowledge, to denying the dynamic nature of language and its link to theorizing, and to classifying whole branches of knowledge—like logic—as language-based and not fact-based. It is thus compatible with a local or relative analytic-synthetic division, but not with a global or absolute division.

An analogy with biology might be useful here (although the limitations of analogies should be kept clearly in mind): Genetically-based traits often stay statistically stable for long periods of time, and this means that inductive generalizations relying on their stability are warranted. Yet it is a central feature of *life* that genetically-based traits are mutable in principle, and those periods in which statistically significant mutations do occur play an exceptionally important role in the development of life. Conceptual connections are similar. Epoch-making changes are rare, yet their significance for human knowledge cannot be overstated. While Grice and Strawson focus on their factual rarity, we focus on their epistemic importance.

(b) Regardless of whether Grice and Strawson are right about Quine's model, we have seen above that an analytic-synthetic distinction that does justice to their first claim is compatible with our model. Yet it is important to realize that this does not eliminate our principled, normative, epistemic objection to a full-fledged analytic-synthetic distinction.

(c) Concerning the threat to other concepts and distinctions, there is nothing in our own rejection of the analytic-synthetic bifurcation (or in any other aspect of our model) to threaten most of the concepts and distinctions Grice and Strawson are concerned about: the concepts of believing or understanding, the distinction between necessary and contingent, sameness and difference of meaning, extension and intension, correct and incorrect translation, meaningfulness and meaninglessness. On the contrary, our model treats these concepts and distinctions as important and significant.¹⁹ We do deny the *absolute* division of units of knowledge into the apriori and the aposteriori and of true statements into those based *exclusively* on reason and those based on fact. But we admit *quasi-apriori* knowledge, and we regard reason as playing a central role in knowledge, indeed a greater and more substantial role than that assigned to it by many adherents of the apriori-aposteriori and analytic-synthetic distinctions.²⁰

¹⁹ Necessity, for example, is central to our account of logic. (See Chapter 10.)

²⁰ There are other criticisms of NAS, or defenses of the analytic-synthetic distinction, that have some relevance to our model, but none, as far as I can see, directly pertains to it. Potentially, Russell's (2008) discussion of *analytic justification* is relevant, but (i) her discussion of this idea is, as she herself indicates,

II. Criticisms of Quine's holism: Grünbaum and Friedman

A. Grünbaum (1960, 1971). Grünbaum's criticism of Quine's holism focuses on the so-called Duhem-Quine thesis (DQ). It is natural to construe DQ as based on the following two principles:

- (i) A principle generalizing the *metalogical* principle:
 - (ML) If \mathfrak{S} is a logically valid inference with a conclusion C and a set of premises P, then: if the cardinality of P is larger than 1, the falsehood of C proves that at least one sentence of P is false, but not that a specific sentence of P is false.²¹
- (ii) A principle based on relational, or network, holism:
 - (RH) The truth, acceptance, and justification conditions of any unit of knowledge are in principle non-trivially connected to those of some other units, usually multiple units.²²

The combination of these principles gives rise to a cluster of theses that can be identified as DQ.

DQ is commonly thought to have originated in Duhem (1906/14) and Quine (1951a). Duhem's version of DQ is usually taken to be the weaker of the two. Duhem limits DQ to empirical science, and within it, to physics, or more generally to mature, highly mathematical science. The conclusions he derives from DQ are therefore limited to those sciences. Duhem's main, or most well-known, conclusion is that *there are no crucial experiments* in physics. That is, there are no experiments that unequivocally decide between competing physical theories or hypotheses. In other words, it is impossible to test a single scientific hypothesis in isolation from other,

quite preliminary, (ii) for the most part, she does not take into account our own reasons for rejecting the distinction, (iii) she does not offer any substantive arguments for her (tentative) view of logic and mathematics as analytic, and (iv) she does not consider the possibility of a new approach to logic and mathematics which can ground them in reality without resorting to either extreme empiricism or radical Platonism, as ours does. (See Chapter 3, Section 4, Chapter 5, Sections 2 and 4, Chapter 8, Section 4, and Chapter 10.) For all these reasons, it is too early to discuss her view in connection with our model. I should indicate, however, that potentially, there are also points of agreement between Russell's approach and ours. Russell regards the relation between language and world as more complex than it is often taken to be, she acknowledges the "temporary profile" of postulations (Russell 2008: 159, see fn. 13), she regards logical statements (e.g., "Snow is white or snow is not white") as *about* the world, her conception of language incorporates something similar to "epistemic freedom", and so on.

²¹ How far (i) generalizes ML is an open question whose answer might vary from one version of DQ to another. For the purpose of the present discussion, this question may be left open. Below I will use "ML" for the generalized principle as well.

²² Here, too, there can be diverse views about the degree and nature of non-trivial interconnections, and this issue, too, can be left aside here. But to guide our understanding, we may think of this connection as taking such forms as a *necessary* or a *sufficient condition* (or some weaker version of these). For example: unit of knowledge U is non-trivially connected to units of knowledge U' if U would not be true/accepted/justified had U' not been true/accepted/justified, or the truth/acceptance/justification of U is (partly) due to the truth/acceptance/justification of U'.

auxiliary, hypotheses, and a single scientific theory apart from other, related, theories. While a cluster of hypotheses (theories) can be tested by an empirical experiment, a single hypothesis (theory) cannot. In Duhem's own words:

[T]he physicist can never subject an isolated hypothesis to experimental test, but only a whole group of hypotheses (Duhem 1906/14: 187).

Quine's (1951a) version of DQ is thought to be stronger—broader and more radical—than Duhem's:²³ Quine applies DQ to all branches of knowledge; he makes the bold claims that any statement whatsoever is open to revision and that any statement can be kept intact in the face of any experimental result; he says that every experiment impacts our entire system of knowledge; he associates DQ with an underdetermination thesis which claims that our system of knowledge is radically underdetermined by experience; and he might be interpreted as saying that all statements/theories are potentially connected to all other statements/theories.²⁴

Grünbaum targets the following Duhemian version of DQ:

(DQ_G) “[T]he *falsifiability* of an isolated empirical hypothesis H... is *unavoidably inconclusive*” (Grünbaum 1960: 116).²⁵

One way to arrive at DQ_G from ML and RH is by treating the “in principle” in RH as an *exceptionless* “always” (or “necessarily”). And this is what Grünbaum, in effect, does. Suppose a hypothesis H predicts (implies, explains) observational result O, but some empirical finding O' is incompatible with O. This is not sufficient to conclusively falsify H, according to the line of reasoning leading to DQ_G. By RH the prediction (implication, explanation) of O by H non-trivially involves a set of assumptions (hypotheses, theories) A, and by ML the finding O' *at most* conclusively refutes the conjunction (or more generally, the combination) of A and H, but not the single hypothesis H.

DQ_G implies that there are no *crucial experiments*, i.e., that no single hypothesis can be decisively refuted by any experiment and, according to Grünbaum, this threatens to undermine a venerable staple of scientific practice.

²³ This specifically relates to Quine's (1951a) version of DQ. In response to Grünbaum's criticisms, Quine (1962) significantly weakened his original version of DQ.

²⁴ This comparison is partly based on the Introduction to Harding (1976).

²⁵ In his (1971) Grünbaum talks about D1 and D2 instead of the thesis that we mark as “DQ_G”, where D1 and D2 are:

- (D1) “No constituent hypothesis H of a wider theory can *ever* be sufficiently isolated from some set or other of auxiliary assumptions so as to be separately falsifiable observationally” (Grünbaum: 261).
- (D2) “For all potential empirical findings O' [incompatible with O], there exists at least one suitably revised set of auxiliary assumptions A' such that the conjunction of H with A' *can be held to be true and explains O'*” (Grünbaum: 261).

For the purpose of our discussion we may treat the conjunction of D1 and D2 as DQ_G.

In a number of papers Grünbaum criticizes DQ_G , and among the general claims and points he makes, the following are especially relevant for understanding the import of his criticism for our own model. Grünbaum admits that:

There are quite trivial senses in which [DQ_G is] uninterestingly true and in which no one would wish to contest [it]. . . . [But in] its non-trivial form, [DQ_G]²⁶ has not been demonstrated (Grünbaum 1971: 262).

Indeed, he claims, the situation is very different when we go beyond the trivial case:

[No] general logical considerations can *guarantee* the deducibility of O' from . . . H and some *non-trivial* revised set A' of . . . auxiliary assumptions which is logically incompatible with A *under the hypothesis* H [S]uch assurance *cannot* be given on general logical grounds at all but . . . the existence of the required set A' needs *separate* and *concrete* demonstration for each particular case. . . . [B]y denying the feasibility of conclusive falsification, the Duhemian schema is a serious *misrepresentation* of the actual logical situation characterizing an important class of cases of falsifiability of a purported explanans [H] (Grünbaum 1960: 118).

Explaining, he says:

[I]t is one thing to maintain . . . that "in every empirical test a certain number of statements of various types is involved" . . . but quite another to conclude in Duhemian fashion . . . that "in principle, it is possible . . . to maintain any particular empirical statement, whatever the data of experience, provided we make appropriate changes in the system of hypotheses which is put to test" (Grünbaum: 116).²⁷

Refining his claim, he says that:

Subject to an important *caveat* to be issued presently, I maintain . . . that there are cases in which we can establish a strong presumption of the falsity of a component hypothesis, although we cannot falsify H in these cases beyond any and all possibility of subsequent rehabilitation (Grünbaum 1971: 283).²⁸

The caveat is:

I emphatically do allow for the following possibility, though not likelihood, in the case of an H which has been falsified in my merely presumptive sense: A daring and innovative scientist who continues to entertain H , albeit as part of a new research program which he envisions as capable of vindicating H , *may* succeed in incorporating H in a theory so subsequently fruitful and well-confirmed as to retroactively alter our assessment of the initial falsification of H . And my *caveat* is that my conception of the falsification of H as establishing the strong presumption of falsity is certainly *not* tantamount to a stultifying injunction to any and all imaginative scientists to cease entertaining H forthwith, whenever such falsification obtains at a given time!

²⁶ Specifically, $D2$.

²⁷ The cited sentences here are from Herbert (1959: 109 and 108).

²⁸ Grünbaum offers an example from physical geometry as a purported proof of this claim. For his example, see (Grünbaum 1960: 119–27). For his response to objections, see (Grünbaum 1971: 278–83). See also appendix to the last paper. For discussions of Grünbaum's example, see e.g., articles in Harding (1976).

Nor is this conception of falsification to be construed as being committed to the historically false assertion that no inductively unwarranted and daring continued espousal of H has ever been crowned with success in the form of subsequent vindication (Grünbaum: 283).

Although our model incorporates a fairly robust version of DQ, Grünbaum's criticism does not apply to it. To see why, let me explain how DQ is construed by our model.

First, we distinguish two aspects of DQ—a positive, constructive aspect, and a negative, destructive aspect. Most of the literature on DQ focuses on the latter, but in our view, the former is at least as central as the latter. DQ, as seen from the perspective of our model, arises from the basic characteristics of the human epistemic situation, like those noted in Chapter 5, Section 1: (i) the complexity of the world, (ii) our limited resources, (iii) our desire to know the world in its full complexity in spite of our limitations, and (iv) our cognitive resourcefulness (which coexists with our cognitive limitations). While the combination of (i), (ii), and (iii) is the source of the negative side of DQ, the combination of (i), (iii), and (iv) is the source of its positive side. As viewed from the perspective of these four factors, what the meta-theoretical principle, ML, and the principle of relational holism, RH, show is that “generally” (in a sense that allows exceptions), given a state of knowledge, K, there are multiple ways of proceeding in the project of knowledge from K on. This applies to states K in which empirical findings O' appear to falsify a scientific hypothesis H by being incompatible with one of its empirical consequences, O.

In such cases, RH says that due to H's interconnection with other units of knowledge (on various levels: physical, logical, etc.), the prediction (implication, explanation) of O by H involves a host of other units of knowledge, A, so that if we take a *broad perspective*,²⁹ we see that O is predicted (implied, explained) not by H alone, but by H in combination with A. It follows from the metatheoretical principle ML that the findings O' can, in principle, be made compatible with H (and adjoined assumptions, hypotheses, theories) by changing one or more of the members of A. What this means, on the positive side, given the other principles of our model, is that no matter how problematic the findings O' are for our current theory T (or hypothesis H), there are in principle multiple routes for overcoming the problem, and that by exercising human ingenuity, we might find some of these. Occasionally, it might even make sense to preserve T (or H) intact by making changes elsewhere in the system. This is the positive, constructive, message of DQ, as it is construed by our model. DQ tells us to continue looking for solutions when a problem arises, to take into account the probable existence of alternative solutions before deciding to adopt a given solution, and in extreme cases, to be ready to go as far as considering a Kuhnian revolution, or a change that involves distant units of knowledge (like logic in the case of physics).

²⁹ Metaphorically, if we move to a relatively high standpoint on Neurath's boat.

But there is a price to pay for this positive aspect of DQ: *uncertainty*. A certain degree of uncertainty is *in principle* involved in any presumed refutation of a given hypothesis H or theory T, both in physics and in other fields. In the literature, this negative side of DQ is often associated (or even identified) with Quine's indeterminacy thesis, which says that empirical theory is underdetermined by observation. This leads to a radically negative construal of DQ: there is no fact of the matter with respect to any hypothesis or theory.³⁰ Many also identify the negative side of DQ with the rejection of *crucial experiments* in science. Replacing our (exception-permitting) "in principle" by (an exceptionless) "always" or "necessarily", they reason that if there are *always* multiple ways of dealing with alleged conflicts between a theory T (a hypothesis H) and a finding O'—including ways that render O' compatible with H—then there is no *decisive test* for any theory, no *crucial experiment* for any physical hypothesis.

There is some force to Grünbaum's objection to this (apparent) result of DQ. How can the general epistemologist or logician tell physicists that no matter how careful they are in testing their, and their colleagues', hypotheses, they can never refute them? This flies in the face of successful scientific practice in fields such as physics. Going further, one might ask: "Given this result of DQ, what is the difference between a carefully designed and executed test or experiment and a sloppy one? Why bother with experiments at all?"

But Grünbaum's objection does not apply to our model. First, as we have noted above, our model regards RH, hence DQ, as holding *in principle*, not *always* in an exceptionless sense. So the claim that there are *no* crucial experiments *whatsoever* does not follow from our version of DQ. Second, DQ, in our model, is but one among a host of principles guiding, and explaining, the pursuit of knowledge, and as such its impact is constrained and limited by other principles. This brings us back to holism and the dynamic structure of our model. Our holism is a *discriminating* holism. It does not say that every unit of knowledge is (significantly) connected to every other unit of knowledge, or that all units are connected to all units to the same extent. This is part of what we meant when we characterized it, in Chapter 2, Section 2, as a structured holism. Interconnections are not *the same in all directions*, and to understand knowledge one has to carefully distinguish between different contexts and different stages (in the development of our system), a principle reflected in the dynamic structure of our model.

Turning back to the basic characteristics of the human epistemic situation, these call for a careful management of our cognitive resources. And one of the resources which is especially relevant to DQ is the human ability to *shift perspectives*—to move

³⁰ What we have said so far about DQ does not require us to accept this radical conclusion. Although I will say a few things that are relevant to this issue below, I prefer to leave the discussion of the indeterminacy thesis to a future occasion (when epistemic freedom is at center stage). For an overview of the current literature on indeterminacy in science, see e.g., Stanford (2009/13).

from one standpoint to another on Neurath's boat. It is quite clear that we cannot make much progress in knowledge if we always assume the broadest perspective possible, and it is equally clear that we will not make much progress if we always assume the narrowest perspective. The project of knowledge requires a balance between a broad array of perspectives, some narrower than others, some wider.³¹ And as far as DQ is concerned, our model recognizes the importance of a standpoint from which some (would be) empirical refutations are significant, compelling, successful—*good enough* to be accepted as *definitive* (during a given stage of the development of science). But this does not conflict with the need, and viability, of a broader perspective, nor with its indispensability both for the development and for our understanding of science. And it is from such a broad perspective that the Duhem-Quine thesis (DQ) gains its force. But DQ itself must be viewed in perspective, one that recognizes the effectiveness of scientific procedures like crucial experiment. By recognizing the importance of both perspectives, the new model is capable of incorporating DQ while being immune to Grünbaum's criticism. Finally, it is worth noting that Grünbaum himself is aware of the positive aspect of DQ, as witnessed by his comments on a "daring and innovative scientist" in the citation above.

To sum up, our response to Grünbaum is that DQ leaves open the question of how one should respond to a particular conflict between a given hypothesis *H* (of a given theory *T*) and an observation *O* that conflicts with its predictions. In principle, there are multiple options for revising our system of knowledge in response to such a conflict, and which revision is warranted in a given case is an open question, to be decided by balancing a host of considerations. It is perfectly compatible with DQ that in the majority of cases the warranted response is the straightforward one of revising *H*, but it is important to recognize other, less frequently warranted but equally important, options. And this our model does. It recognizes that testing individual hypotheses by means of experiment is an important element of scientific method, and it also recognizes that an occasional divergence from this method is fundamental both to the actual development of science and to our understanding of its development. Grünbaum's criticism is directed at a different, more rigid construal of DQ than the one offered by our model.

Let us conclude with another criticism of Quine's holism, due to Friedman.

B. *Friedman (1997, 2001)*. Friedman is highly critical of two aspects of Quine's holism: (i) its *symmetric, egalitarian* treatment of all units of knowledge, and (ii) the *naturalistic* character of his holism. Friedman describes this holism as follows:

Quine's epistemological holism pictures our total system of science as a vast web or conjunction of beliefs that face the "tribunal of experience" as a corporate body. Quine grants that

³¹ We will discuss a similar balance in connection with truth in Chapter 7, Section 2.

some beliefs, such as those of logic and mathematics, for example, are relatively central, whereas others, such as those of biology, say, are relatively peripheral. But this means **only** that the former beliefs are **less likely** to be revised in case of a “recalcitrant experience” at the periphery, whereas the latter are **more likely** to be revised. **Otherwise, from an epistemological point of view, there is simply no relevant distinction to be made here...**

[E]mpirical evidence—either for or against—spreads over all the elements of the vast conjunction that is our total system of science, wherein all elements whatsoever equally face the “tribunal of experience”. And it is in this precise sense, for Quine, that **all beliefs whatsoever, including those of logic and mathematics, are equally empirical** (Friedman 2001: 33–5).

But this kind of holism, Friedman argues, is not characteristic of actual scientific practice; nor is it normatively sound:

[T]his beguiling form of epistemological holism [**does not**] **do justice** to the revolutionary developments within both mathematics and natural science. [Ibid. Friedman: 35]

[N]aturalistic epistemological holism... is **entirely incapable** of providing an adequate philosophical perspective on [the exact] sciences (Friedman: xii).

[Quine’s naturalistic holism] has reached the **end of its useful life** (Friedman 1997: 7).

One root of the problem, according to Friedman, is the *undiscerning, symmetrical* treatment of all constituents of knowledge by Quine’s holism, overlooking the substantial differences between them. Consider Newtonian physics:

[T]here are fundamental asymmetries in the way in which the different elements of [the] Newtonian synthesis actually function.... The combination of calculus plus the laws of motions is not happily viewed... as a conjunction of elements symmetrically contributing to a single total result (Friedman 2001: 35–6).

[A]dvanced theories in mathematical physics, such as Newtonian mechanics and Einsteinian relativity theory, should be viewed as consisting of two **asymmetrically functioning parts**: a properly empirical part containing laws such as universal gravitation, Maxwell’s equation of electromagnetism, or Einstein’s equations for the gravitational field; and a constitutively a priori part containing both the relevant mathematical principles used in formulating the theory (Euclidean geometry, the geometry of Minkowski space-time, the Riemannian theory of manifolds) and certain particularly fundamental physical principles (the Newtonian laws of motion, the light principle, the equivalence principle). (Friedman: 71).

It is important to recognize, Friedman says, that the evidence appropriate for different units of knowledge is *not* the same (or of the same kind):

It is a profound mistake of Quinean holism... to view [the] procedure [of experimentally testing the bona fide physical laws which presuppose certain mathematical and general scientific laws] **as empirically testing the other two parts in the same way**. For... it is clear that the mathematical part of our theories, considered independently of the empirical application in question, is in no way empirically tested by such a procedure (Friedman: 80).

Similarly, Friedman argues, Quine's view that what distinguishes the mathematical part of empirical science from other parts—namely, psychological features like a greater degree of entrenchment—is incorrect, both historically and normatively:

When Newton formulated his theory of gravitation, ... the **mathematical part** of his theory, the new calculus, was **still quite controversial**. ... Similarly, in the case of Einstein's general theory of relativity, **neither the mathematical theory of manifolds nor the principle of equivalence was a well-entrenched part of main-stream mathematics or mathematical physics**. ... [S]ince we are dealing with deep conceptual revolutions in both mathematics and mathematical physics in both cases, **entrenchment and relative resistance to revision are not appropriate distinguishing features at all** (Friedman: 39–40).

Most significantly, Friedman rejects Quine's holism due to its *naturalistic* character—in effect, its radical bias toward the *empirical*. Referring to Quinean holism as “philosophical naturalism”, he says:

This tendency of thought, which I will call “philosophical naturalism”, ... is characterized by two main ideas. The first is the rejection of any special status for types of knowledge traditionally thought to be a priori—knowledge in logic and mathematics, for example—in that **all knowledge whatsoever is now conceived as having fundamentally the same status as that found in the empirical natural sciences**. ... The second main idea ... is the view that philosophy, as a discipline, is also best understood as simply one more part ... of **empirical natural science** ... [reflected in Quine's claim that] “epistemology, or something like it, simply falls into place as a chapter of psychology and hence of natural science” (Friedman 1997: 7).

Friedman rejects both ideas.

Neither of Friedman's criticisms of Quine's holism applies to the type of holism advocated by our model:

(i) *Differentiation Between Different Branches of Knowledge*. Our holism does not treat all units of knowledge symmetrically, that is, in the same manner. It is rooted in (a) the search for a (foundational) alternative to traditional foundationalism, and (b) recognition that our basic epistemic situation requires, and gives rise to, a rich, complex, and open-ended structure of cognitive routes to reality. There is nothing in either of these roots to require blind symmetry—the view that all units of knowledge are (significantly) connected to *all* other units of knowledge, or that they are all of the *same* kind, or that they are connected to all units (that they are connected to) in the *same* way, to the *same* extent, in *all* contexts, or at *all* stages of their development. On the contrary. This is an aspect of our holism that was emphasized already in Chapter 2, Section 2. Our holism is a *structural* holism, a holism that *recognizes* and *distinguishes* different positions in the structure of knowledge (on multiple dimensions) and differentiates disciplines according to the positions they occupy as well as the ways they make these positions their own. Take abstract and experimental units of knowledge. Our model distinguishes between them with respect to what features of reality they study (abstract or physical), and accordingly, with respect to what evidence and justification procedures they can, do, and ought to use.

(ii) *Extreme Empiricism and Assimilation of Philosophy to Empirical Science*.³²

Holism comes in different forms, as we have seen in Chapter 2, Section 2. In one of its forms, it is a coherentist methodology, in another a foundational methodology (an alternative to the rigid foundationalist conception of the way knowledge is grounded in reality). Similarly, holism can be conjoined with radical empiricism, but can also be conjoined with a more inclusive view of knowledge, one that views intellect, no less than sensory perception, as a major source of knowledge and emphasizes the cooperation between the two sources. Our own holism is of the second kind. By recognizing the substantial difference between these two sources, it affirms their independence; by recognizing their interconnections, it licenses their cooperation. By affirming their independence it denies that either is reducible to the other; by emphasizing their cooperation it rejects the view that (for the most part) they operate in complete isolation of each other. Either way, it rejects both radical apriorism and radical empiricism. As a result, Friedman's view that mathematics is both independent of science and plays a central role in science is fully compatible with our brand of holism, and indeed it is seconded by our model. The same holds for his view of philosophy. Philosophy, in our model, is both independent of other sciences and engaged in fruitful cooperation (and mutual fertilization) with them, as emphasized in Chapter 3, Section 5.³³

Furthermore, Friedman's *positive* conception of the *structure* of scientific knowledge is also compatible with our holism. Friedman describes this conception as follows:

Let me suggest . . . [a] picture of a dynamical system of beliefs, concepts, and principles that can be analyzed . . . into three main components: an evolving system of empirical natural scientific concepts and principles, an evolving system of mathematical concepts and principles which frame those of empirical natural science and make their rigorous formulation and precise experimental testing possible, and an evolving system of philosophical concepts and principles which serve, especially in periods of conceptual revolution, as a source of suggestions and guidance in choosing one scientific framework rather than another. All of these systems are in continual dynamic evolution. . . . Yet we can nonetheless clearly distinguish the radically different functions, levels, and roles of the differing component systems. In particular, although the three component systems are certainly in perpetual interaction, they nonetheless evolve according to their own characteristic dynamic (Friedman 1997: 19).

Both this tripartite structure and its dynamic evolution can be seamlessly incorporated in our model.

³² Throughout the essay I prefer to talk about the radically empirical character of Quine's holism rather than its naturalistic character. The reason is that naturalism has many sides to it and is understood in a variety of ways within the philosophical community, while radical empiricism is better demarcated and there is a greater unity in the way it is understood.

³³ I should perhaps add that our model does not characterize the "center" (of our system of knowledge) in terms of "entrenchment". On the contrary, disciplines traditionally located in the center are, and ought to be, located in the periphery as well, where no matter how entrenched they are, they are subject to critical evaluation in light of their veridicality, namely, being true to reality.

Friedman, in fact, explicitly endorses some features of holism:

[I]t is indeed the case that **no concept or principle is forever immune to revision** (Friedman: 19).
[T]he idea of a special...role for the mathematical disciplines in our natural scientific knowledge...**has nothing to do** with a jejune obsession with **epistemic certainty, unshakable foundations, or absolute unrevisability** (Friedman: 15).

We endorse the same features.

Friedman's description of his motivation for viewing mathematics as playing a "special role" in science is also shared by our model:

[My view] is motivated throughout by an appreciation of the **manifold possibilities for development, growth, and radical transformation** in both pure mathematics and mathematical natural science—and by appreciation, above all, of the **striking and unexpected ways** in which these two types of developments can **influence and even merge with one another** in the course of revolutionary conceptual changes such as those exemplified in the theory of relativity (Friedman: 15).

These observations, indeed, are among those underlying our endorsement of holism.

Similarly, Friedman's conception of knowledge as highly dynamic is shared by our model:

In place of the Quinean figure of an holistically conceived web of belief,...I would like to suggest an alternative picture of a **thoroughly dynamical yet nonetheless stratified or differentiated** system of knowledge (Friedman 2001: 45).

Ironically, perhaps, it is the holistic nature of our model that makes its dynamic character possible. It is the structural nature of this holism that makes its highly structured—"stratified or differentiated"—conception of knowledge possible.

Finally, Friedman's view that something like apriority—yet different from the traditional conception of apriority—is essential for knowledge, is also shared by our model (as we have seen in Chapter 5, Section 3).

Summing up the different elements brought together by our model of knowledge: it exemplifies the principles of epistemic friction and freedom with which we opened this essay; it incorporates the holistic foundational methodology developed in Chapter 2, Section 2; it is based on the nine building blocks delineated in Chapter 3, Section 1; it takes steps toward realizing the promise of Quine's model while avoiding its problems; it is thoroughly dynamic, requiring disciplines lying at the "center" of Quine's model to shift, in some contexts and times, to the "periphery", where they are subject to strong veridicality norms, including grounding-in-reality; it rejects the traditional analytic-synthetic thesis (though not on the same grounds as those emphasized by Quine in most of his writings); it rejects extreme empiricism, but not the centrality of empirical tests for science; it advances a conception of reality that goes beyond the traditional doctrines of nominalism and Platonism, affirming

abstract features of reality yet locating them in the same reality as physical objects (as features of these objects, their properties and relations); it assigns a central role to intellect in knowledge—including discovery and justification—without denying the concurrent centrality of sensory resources; it questions the traditional emphasis on intellect's "pure" activity, or the ways it operates in *complete isolation* from experience ("apriorism"); it upholds a robust yet flexible form of realism—*basic realism*; and it is immune to many of the influential criticisms of Quine's model, including criticisms directed at his treatment of analyticity and his particular brand of holism.

In Part III, we will extend our theory to the domain of truth. Truth enters into our theory in two ways. First, we view *truth itself* as a universal standard of knowledge. Truth is associated with the universal friction requirements of veridicality and grounding-in-reality. The universality of these requirements is reflected, in our model, by the fact that both normatively and descriptively, *all* branches of knowledge lie in the periphery (and not just in the center or the intermediate zones). The relevant sense of periphery, here, is contact with reality, and reality itself is viewed as both the target of every field of knowledge (different aspects of reality being the targets of different fields) and the focus of the veridicality norms governing it. This by itself does not determine a unique theory of truth, but it sets significant constraints on the development of such a theory. The fact that as far as veridicality is concerned, fields traditionally viewed as occupying Quine's center (and which in Quine's model have no more than indirect connections with the periphery) are viewed, in the present model, as lying squarely in the periphery, makes it especially important to develop a robust theory of truth that applies to those branches. And the combination of a robust norm of truth and a view of reality that goes beyond both Platonism and empiricism, introduces a challenge that cannot be swept under a deflationist, or a minimalist, approach to truth. This leads to the second way in which truth enters into our theory, namely through the character of the theory of truth as a philosophical theory.

The *theory of truth*, as a *philosophical* theory, is at the center of the contemporary debate between deflationists (minimalists, quietists) and substantivists on the character of philosophy. By analyzing the difficulties involved in developing a substantive theory of truth we will be able both to offer a new perspective on this controversy and to develop guidelines for overcoming the underlying difficulties. Substantiveness is one of the universal friction requirements governing our model, and confronting the challenges it poses is a step toward demonstrating the viability of a substantivist approach to philosophical theorizing and establishing philosophy as a genuine field of knowledge within the model, alongside physics, mathematics, and other fields.

Accordingly, our investigation of truth will focus both on the structure of the *theory of truth* and on the structure of *truth itself*. In Chapter 7, we will confront some of the methodological difficulties involved in constructing a *substantivist* (as opposed to a *deflationist*) theory of truth and devise a strategy for overcoming these difficulties. In Chapter 8, we will use this strategy to investigate the *principles* of truth itself. Chapter 7 will focus on the *unity and disunity* of truth; Chapter 8 on the *connection between mind and reality* constitutive of truth.

PART III

The Structure of Truth

A Substantivist Theory of Truth

7.1 A Substantivist (as Opposed to Deflationist) Methodology

One of the main friction requirements for all theories in our corpus of knowledge is substantiveness. Today, this requirement is especially significant in the field of truth due to philosophers' attraction to the deflationist (anti-substantivist) approach. In Chapter 1, Section 2 we introduced the substantiveness requirement as concerned with "high standards of discovery, explanation, justification, informativeness, depth, theoretical significance, rigor, systematicity, intellectual interest, and the like". This list was intended not as a definition, a necessary and/or sufficient condition, or a check list, but rather as a broad and open-ended indication of the kind of features we value in a substantive theory. In considering philosophical theories, we may add further desiderata; for example, theorizing about a significant, rich, and fundamental subject matter, and theorizing constructively. These, like all the above attributes, are understood in their everyday, common-sense meaning.

Our conception of a substantive philosophical theory resonates with those of Dummett (1976/91) and Goldman (1979). Dummett opens his 1976/91 lectures by saying:¹

The layman... expects philosophers to answer **deep questions of great import for an understanding of the world**.... And the layman is **quite right**.... Yet he finds most writings by philosophers of the analytical school disconcertingly remote from these concerns.... The complaint... is understandable;... analytical philosophy passed, comparatively recently, through a destructive phase.... During that phase, it appeared as though demolition was the principal legitimate task of philosophy. Now most of us believe once more that philosophy has a **constructive task**; but, so thoroughly was the demolition accomplished, that the rebuilding is of necessity slow (Dummett 1976/91: 1).

And Goldman describes the goal of his theory of knowledge as:

The aim... is to sketch a theory of justified belief. What I have in mind is an **explanatory theory**.... I want a set of **substantive conditions**.... I seek an explanatory theory, i.e., one that **clarifies the underlying source** of justificational status.... Its conditions must... be

¹ For **boldface** within citations, see fn. 3, Chapter 1.

appropriately **deep** or **revelatory**. . . . A theory of justified belief of the kind I seek . . . must be couched at a suitably **deep** . . . **level** (Goldman 1979: 1–2).

The substantivist approach, in the sense intended here, contrasts with the *deflationist* approach to truth, or certain aspects of this approach.² “Deflationism” is a broad and diffused category, used in different ways in different contexts. Here I focus on, and indeed limit myself to, that aspect of deflationism that claims that the notion of truth is *trivial* and advocates a *shallow* theory of truth. This aspect of deflationism is reflected in views like the following:

[T]ruth is **entirely captured** by the . . . **triviality** . . . that each proposition specifies its own condition for being true (e.g. the proposition *that snow is white* is true if and only if *snow is white*), . . . so that in fact nothing could be more **mundane** and **less puzzling** than the concept of truth (Horwich 1990/8: xi).³

[It is a] misconception . . . [to think] that truth *has* some **hidden structure** awaiting our **discovery**. . . . Unlike most other predicates, ‘is true’ . . . should not be expected to participate in some **deep theory** of that to which it refers—a theory that goes beyond a specification of what the word means (Horwich: 2).⁴

Substantivists (advocates of a substantive theory of truth) differ from deflationists on multiple points: Where deflationists say that “truth is entirely captured by the . . . triviality . . . that each proposition specifies its own condition for being true”, substantivists say that it is far from being fully captured by this triviality; where deflationists say that “the truth predicate exists solely for the sake of a certain logical need”⁵ (Horwich: 2), substantivists say that it exists for other needs as well; where deflationists say that truth is not a deep notion, substantivists say it is; and where deflationists say that a theory of truth cannot be, or need not be, genuinely explanatory, substantivists say it can and should be.⁶ Substantivists accept the view that truth is (initially) mysterious, but believe its mysteries can be solved rationally. This notion of a substantive theory is close to what some call an “inflationary theory”⁷, but “substantive”, in its everyday usage, better captures the ordinary, common-sense considerations that motivate us. Like Dummett we believe that philosophy should take on important “questions . . . for . . . understanding . . . the world”, and like the early

² Deflationism, as it is intended here, encompasses also some versions of *minimalism* and *quietism*.

³ The original order of sentences was changed. The original text reads: “Perhaps the only points about truth on which most people could agree are, first, that each proposition specifies its own condition for being true (e.g. the proposition *that snow is white* is true if and only if *snow is white*), and, second, . . . I shall show that truth is entirely captured by the initial triviality, so that in fact nothing could be more mundane and less puzzling than the concept of truth.”

⁴ Horwich is a so-called “minimalist”, but his minimalism falls under the category of “deflationism” as it is used here.

⁵ I.e., the need for indirect reference to, and generalization over, sentences or propositions.

⁶ Even those deflationists who do not eschew explanation altogether normally have in mind a far more limited explanation than what substantivists have in mind. For an insightful critique of deflationists’ conception of explanation, see C. D. Wright (2015).

⁷ See e.g., Field (1994).

Field we think that “we’d be crazy to give . . . up” in philosophy “a methodology [of substantive theorizing] that has proved extremely fruitful in science” (Field 1972: 363).⁸

The attempts to construct a substantive theory of truth, however, have come upon great difficulties, and many philosophers have given up hope of ever producing such a theory. Field, for example, has renounced his plan for a substantive theory of truth (based on a causal account of reference), and today he, along with many adherents of his original plan, champion deflationism.⁹ Not all contemporary philosophers, however, are satisfied with the prevalent trend. Putnam (1981, 1987), Devitt (1984/91), Davidson (1990, 1996), Wright (1992, 1999), Gupta (1993a, 1993b), David (1994), Lynch (1998, 2004a, 2009), Sher (1998–9, 2004¹⁰), C. D. Wright (2015)¹¹, and others have dissented, to a greater or lesser extent, from deflationism.¹²

In this essay I will approach deflationism through the question: “What underlying reasons might lead philosophers to adopt this position, other than the usual refrain that otherwise truth is a mystery or the observation that there are uses of truth that do not require a substantive conception of truth?”¹³ The problem with the “otherwise truth is a mystery” refrain is that it is based on a negative existential assumption, the assumption that there is no way to provide a substantive explanation of truth, an explanation that would remove the “mystery”. Negative existential claims, however—in particular, methodological negative existential claims—are very difficult to establish. How can one establish “the fact” that it is impossible to come up with a substantive account of truth? In particular, how can one rule out the possibility that it is because of the way we customarily think about truth that we find it so difficult to come up with a substantive account of it, i.e., that it is because we overlook something fundamental about truth that we have difficulties in solving its “mystery”? Furthermore, starting with a mystery is not a bad thing either in philosophy or in other fields. Often, mysteries give us the impetus to investigate a given subject matter thoroughly and in depth, search for a solution to an important question, or look for an explanation of something intriguing and significant. Similarly, it does not follow

⁸ Field treats the theory of truth as part of linguistics, but what he says is directed at whatever discipline the theory of truth belongs to, i.e., on our demarcation, philosophy. The substantive theorizing he has in mind is reductionist, but it is reasonable to presume that he intends his point to extend to substantive theorizing more generally.

⁹ See e.g., Field (1986, 1994). It should be noted, though, that Field himself has always maintained a certain degree of ambivalence toward deflationism.

¹⁰ Incorporated here.

¹¹ To distinguish between Crispin Wright and Cory Wright I use “Wright” for the former and “C. D. Wright” for the latter.

¹² “To a greater or lesser extent” is intended to indicate that there are significant variations in the degree to which the above authors support a substantivist approach to truth (in our sense), and that some express doubts about the viability of such an approach. Still, their overall view of theorizing about truth, as I understand it, is more substantivist than deflationist (in the sense of “deflationist” relevant here, which involves advocating, or at least sanctioning, triviality).

¹³ Both the refrain and the observation are implicit in Horwich (1990/8: 1–3).

from the fact that we can do something useful with a shallow notion of truth that we do not need a substantive notion of truth to do other useful and/or important things.

In particular, I will ask: “Are there any deep methodological challenges to a substantive theory of truth that might discourage philosophers from pursuing such a theory?”. And if the answer is positive, my next step will be to ask whether these challenges arise in other areas besides truth and how they are dealt with in these areas. This strategy will enable me to place the debate on truth in a new perspective, one that is broader and, possibly, more fruitful than current perspectives.

Following this strategy, I will identify a pair of methodological challenges that face our system of knowledge in many areas, including science and philosophy. These two challenges are complementary: the challenge of *disunity* and the challenge of *unity*. In the case of truth, these are the challenges of recognizing the diversity, complexity, and multidimensionality of truth, and of figuring out the structure of its unifying principles. Once we understand these challenges and devise guidelines for meeting them, the road will be open for a substantivist theory of truth.

7.2 The Unity and Disunity of Truth: Challenges and Strategies

The disunity challenge, as it applies to truth, concerns both truth itself and the theory of truth. The two parts of this challenge are interconnected. If truth is diverse in the extreme, there is no basis for a unified substantive theory of truth; and if we, humans, lack, and are in principle incapable of developing, resources for detecting the unified substance of truth, then, as far as we are concerned, it does not have such substance. Below I will distinguish two disunity challenges to substantive theorizing about truth, the one radical, the other moderate. I will argue that the radical disunity challenge is based on unreasonable assumptions, but the moderate challenge is genuine and solving it is a key to a unified substantive theory of truth.

I. Disunity challenges

A RADICAL DISUNITY CHALLENGE

A well-known argument against the feasibility of a substantive theory of truth says that since every thought (proposition, belief, cognition, judgment, sentence, statement, theory)¹⁴ has its own unique truth condition, a general and substantive account of truth is impossible. I will call this argument “the radical disunity argument”. A contemporary formulation of this argument is:

[C]ompare ‘is true’ . . . with a genuine target of philosophical analysis. . . . We know *individually* what makes [‘is true’] applicable to the judgements or sentences of an understood language. ‘Penguins waddle’ is a sentence true, in English, if and only if penguins waddle. It is true that

¹⁴ For the purpose of this essay, there is no need to sort out the different truth bearers.

snow is white if and only if snow is white. The reason the first sentence deserves the predicate is that penguins waddle, and the reason why the judgement that snow is white deserves the predicate is that snow *is* white. But these reasons are entirely different. There is no single account, or even little family of accounts, in virtue of which each deserves the predicate, for deciding whether penguins waddle has nothing much in common with deciding whether snow is white. There are *as* many different things to do, to decide whether the predicate applies, as there are judgements to make. So how *can* there be a unified, common account of the “property” which these quite different decision procedures supposedly determine? (Blackburn 1984: 230).¹⁵

One way of understanding the radical disunity argument is in terms of a “criterion” of truth.

(a) *The Myth of the Criterion*. An early, theoretical version of the radical disunity argument is found in Kant’s introduction to “Transcendental Logic”, *Critique of Pure Reason*. Inquiring whether it is possible to go beyond a minimalist characterization (literally, name-clarification, *Namenerklärung*) of truth as agreement of a cognition with its object and provide a *criterion* that determines the precise conditions under which each cognition is true, Kant reasons:

If truth consists in the agreement of cognition with its object, that object must thereby be distinguished from other objects; for cognition is false, if it does not agree with the object to which it is related, even although it contains something which may be valid of other objects. Now a general criterion of truth must be such as would be valid in each and every instance of cognition, however their objects may vary. It is obvious however that such a criterion [being general] cannot take account of the [varying] content of cognition (relation to its [specific] object). But since truth concerns just this very content, it is quite impossible, and indeed absurd, to ask for a general test of the truth of such content. **A sufficient and at the same time general criterion of truth cannot possibly be given. . . . Such a criterion would by its very nature be self contradictory** (Kant 1781/7: A58–9/B83).¹⁶

A somewhat different formulation appears in Kant’s *Lectures on Logic*:

A universal material criterion of truth is not possible; it is even self-contradictory. For as a *universal* criterion, valid for all objects in general, it would have to abstract fully from all difference among objects, and yet at the same time, as a material criterion, it would have to deal with just this difference, in order to be able to determine whether a cognition agrees with just that object to which it is related and not just with any object in general, in which case nothing would really be said. . . . [I]t is absurd to demand a universal material criterion of truth, which

¹⁵ This argument has also been viewed as an argument against the existence of a *property* of truth, but here I am interested in it as an argument against the possibility of a substantive account of truth. Blackburn, it should be noted, does not present this argument as one he endorses.

¹⁶ (i) I am taking the liberty of changing Kemp Smith’s translation of “*Erkenntnis*” from “knowledge” to “cognition”.

(ii) Here, the bracketed clarifications are due to Kemp Smith.

should abstract and at the same time not abstract from all difference among objects (Kant 1770s–1800: 558).

Considering Kant's argument in the context of a substantive theory of truth, we may interpret it as saying that, to the extent that a *general substantive theory of truth* has to provide a *general criterion* for truth (in Kant's sense), it would face an irresolvable conflict between generality and particularity.¹⁷ Concisely, we may sum up this version of the radical disunity argument as follows:

1. Truth consists in the particular agreement of a thought with its unique object.
2. A general theory of truth must, in order to be general, abstract from the particularity of this relation.
3. But a substantive theory of truth cannot (if it is to be substantive) abstract from its particularity. Hence:
4. A general and substantive theory of truth is impossible.

Now, on some level, this argument is persuasive. Not only is it formally valid, but its conclusion is, in some sense, correct. There is an intuitive sense in which to understand why a certain sentence is true (or false), we have to turn to the field of knowledge it belongs to. To understand why (in virtue of what) the Löwenheim-Skolem theorem is true we have to turn to model theory; to understand why creationism is false we have to turn to modern physics and biology as well as to theology and scientific methodology; to understand why "Picasso revolutionized modern art" is true we have to turn to art history, and so on. It is absurd to think that the theorist of truth could come up with a general and substantive criterion (or necessary and sufficient condition) that would determine, all by itself, the truth value of each and every truth bearer. But in a deeper and more important sense the radical disunity argument is unsound.

This argument assumes an altogether unreasonable conception of a substantive theory of truth, namely, as consisting of, or offering, a Kantian criterion of truth.¹⁸ In fact, there is no reason that a substantive theory of truth, a theory aiming at a philosophical explanation of truth, should be interested in, or be required to provide, such a criterion. Achieving a genuine understanding of truth does not mean detecting all the minute differences between any distinct truths, or determining what exactly has to be done in order to find out whether such sentences as "Penguins waddle" and "Snow is white" are true. There is a whole array of intermediate projects between the minimalist (deflationist) project of name-clarification of truth and the maximalist project of providing a Kantian criterion of truth. A substantive theory of truth should focus on one or more of the intermediate projects.

¹⁷ This Kantian argument received little attention by most Kant scholars. Among those who did discuss it are Prauss (1969), Van Cleve (1999, Chapter 13), Hanna (2000), Rosenkoetter (2009), and Vanzo (2010).

¹⁸ By a "Kantian criterion of truth" I mean, in this essay, what Kant, in the above citations, referred to as a "criterion of truth".

It is not clear whether Kant himself posed the false dilemma of “either our account of truth is (what we call today) deflationist, or it is committed to a criterion of truth”. On the one hand, in presenting his argument he does appear to contrast the mere “name clarification” of truth (his conception of a deflationist theory) with a *criterion* of truth. On the other hand, elsewhere he repeatedly affirms the existence of substantive accounts of various facets of truth and makes substantive claims about truth. For example, he characterizes Transcendental Analytic as “a logic of truth” (Kant 1781/7: A62/B87 and elsewhere), saying it sets negative but not narrowly logical conditions on the possibility of truth; he says that “certain concepts...render possible...all...truth...of empirical cognition” (Kant: A125); he claims that “transcendental truth...precedes all empirical truth and makes it possible” (Kant: A146/B185); he suggests that an account of “the formal conditions of empirical truth” is possible (Kant: A191/B236); he identifies conditions under which “cognition *a priori*...can possess truth” (Kant: A157/B197); he proclaims that the principle of causality is a “condition of...empirical truth” (Kant: A202/B247); he argues that the possibility of “experience” is a necessary condition for [concepts related to objects, i.e., bearing] “truth” (Kant: A489/B517); he implies that a “sufficient mark of empirical truth” is possible (Kant: A651/B679)¹⁹; and, using “criterion” in a more reasonable way than in the introduction to “Transcendental Logic”, he contends that his own theory, unlike Berkeley’s, is capable of providing a “certain [i.e., sure] criterion for distinguishing truth from illusion” (Kant 1783: 124). In all these passages Kant says things about truth that go beyond mere “name clarification” but fall short of a “criterion” in the strong sense of a single condition that determines all by itself, with respect to each and every judgment, whether it is true or false.²⁰

Be that as it may, the dilemma raised by the radical disunity argument is a false dilemma. The radical disunity argument rejects the possibility of a substantive theory of truth by arguing against an *absurd* conception of such a theory.

Another way to see the absurdity of the radical disunity argument is to note that it is applicable, in principle, to other philosophical theories as well, for example, the theory of knowledge. But as far as I know, no one—least of all Kant—used it in order to reject the possibility of a substantive theory of knowledge. Indeed, the disunity argument against the theory of knowledge could be expressed in Kant’s own words (where expressions relating to truth are replaced by expressions related to knowledge):

A universal material criterion of knowledge is not possible; it is even self-contradictory. For as a *universal* criterion, valid for all objects [of *knowledge*] in general, it would have to abstract

¹⁹ Here I am using the Guyer and Wood translation, since they, like Kant, use “mark” (“*Merkmal*”), whereas Kemp Smith uses “criterion” (“*Kriterium*”), which is less accurate, especially in the present context.

²⁰ Hanna (2000, 2004/13), for example, suggests that although Kant did not provide a general criterion of truth in the strong sense of the Introduction, he provided “special criteria of truth for each of the basic classes of judgments: analytic judgments, synthetic a posteriori (or empirical) judgments, and synthetic a priori judgments” (Hanna 2004/13: 9). I would replace “special criteria” by “special *principles*”, since at least with respect to synthetic judgments, these “special criteria” do not purport to tell, all by themselves and in advance, which individual judgment of a given type is true and which is false.

fully from all difference among objects, and yet at the same time, as a material criterion, it would have to deal with just this difference, in order to be able to determine whether a true cognition is *justified* with respect to just that object to which it is related and not just with any object in general, in which case nothing would really be *known*. It is absurd to demand a universal material criterion of *knowledge*, which should abstract and at the same time not abstract from all difference among its objects.

This argument says that knowledge consists in a particular (true-and-justified) belief about a particular object. A general theory of knowledge must, in order to be general, abstract from the particularity of any specific objects of knowledge. But a substantive theory of knowledge cannot (if it is to be substantive) abstract from this particularity. Hence, a general and substantive theory of knowledge is impossible.

Yet in spite of the similarity between truth and knowledge with respect to the radical disunity argument, Kant did not use this argument to demonstrate the unviability of a substantive theory of knowledge. Indeed, Kant's theory of knowledge in the *Critique of Pure Reason* is a *paradigm* of a substantive philosophical theory. The point is that no one in her right mind would require a substantive theory of knowledge to provide a full and detailed criterion of knowledge, a criterion determining, *all by itself*, with respect to each and every judgment, whether it should be included in our corpus of knowledge. Why, then, would anyone think that a substantive theory of truth is required to provide such a criterion? Just as the theory of knowledge is not assigned the task of a (Kantian) criterion, so the theory of truth should not be assigned it. One might argue that the theory of truth is different because truth is different: the truth of a thought *is* dependent on its specific content and object. But doesn't the same hold for knowledge? (Aren't the knowledge conditions of "Penguins waddle" and "Snow is white" just as particular and just as diverse as their truth conditions?) If you start from an unreasonable conception of the task of the theory of truth, it is no wonder that you arrive at a skeptical conclusion. If you think of the theory of truth in "either-or" terms—the theory of truth is either a criterial theory or a deflationist theory—you have no choice but to opt for the latter. But there are many other options.

The radical disunity argument is thus best interpreted as a *reductio ad absurdum* of both the *criterial* view of a substantive theory of truth and the view that truth itself is *radically particular* (each sentence has its own type of truth conditions). Truth is no more particular than knowledge, and a substantive theory of truth, like a substantive theory of knowledge, must abstract from, i.e., overlook, some of the differences between its objects. The choice is not between a deflationist theory that pays no attention to any differences between truths and a criterion of truth that pays attention to all their minute differences; the choice is between these two and a substantive theory of the *major* principles of truth and the *major* differences between truths.

Another way of understanding the radical disunity argument is as invoking a *common denominator* of all truths.

(b) *The Myth of the Common Denominator.* We have seen that the criterial argument against a substantive theory of truth is based on a misconception of such a theory. But there is a more common misconception of a substantive theory of truth that leads to a similar conclusion. We may view it as another version of the radical disunity argument. According to this version—which I will call the “common denominator” version—the task of a substantive theory of truth is to provide a *definition*, or formulate a *necessary and sufficient condition*, of truth (or of “P is true”), one that captures *the common denominator of all truths*. Such a definition (condition) is supposed to account for the *full “substance”* of truth, or delineate *all the “determinants”* (*substantive* factors, constituents, explanans) of truth. This version of the radical disunity argument says that there is no substantive common denominator to all truths, that truth is too diverse to have such a denominator.

To understand the grounds for this version of the argument, consider two truths of the same syntactic structure but altogether different subject matters, say, “ $2 < 10$ ” and “John loves Mary” (assuming it to be true). In spite of their syntactic similarity, the substantive truth conditions of these two sentences appear to have nothing in common. The truth of “ $2 < 10$ ” is a matter of such things as the position of 2 and 10 in the natural number series, the non-existence of surjective (“onto”) mappings from sets with two elements to sets with ten elements, and so forth. The truth of “John loves Mary”, in contrast, is a matter of John’s feelings towards Mary, John’s behavior regarding Mary, John’s brain states involving a representation of Mary, and so on. The truth conditions of these two sentences are altogether different, and an informative theory of truth must account for these differences. But a general theory of truth must abstract from these differences. Ergo: a contradiction. *A substantive theory of truth is not general; a general theory of truth is not substantive.* To invoke a general principle which says that a sentence of the form “*Rab*” is true if and only if (iff) the referent of “*a*” stands to the referent of “*b*” in the relation referred to by “*R*” is to gloss over pertinent differences. To assign each sentence its own unique truth condition is to preclude a general explanation. A general account of truth does not *explain* what makes sentences true, and an explanatory account of truth is not *general*. In short, a theory of truth which is both general and substantive is impossible.

To see the unreasonableness of the common-denominator conception of a substantive theory of truth, let us think of truth as a property²¹, and let us call a property *P* *perfectly singular* (or *perfectly particular*) if for every *a* and *b* in its domain, if $a \neq b$, then the conditions under which *a* possesses *P* are essentially different from the conditions under which *b* possesses *P*. There is no doubt that truth is singular to

²¹ In Chapter 8 I will argue that it is fruitful to view truth primarily as a standard and only secondarily as a property, but for the purposes of the present chapter we may follow the more traditional approach and regard it as a property.

some degree: the truth conditions of “Penguins waddle” and “Snow is white” are different from each other, and the truth conditions of “ $2 < 10$ ” and “John loves Mary” are *even more* different one from the other. But the radical disunity argument involves the assumption that truth is *perfectly*, or at least *very highly* singular, that is, that the truth of *any* two sentences is based on *altogether* different principles. This is clearly incorrect. While the truth conditions of “ $2 < 10$ ” and “John loves Mary” are quite different from each other, there are numerous sentences whose truth conditions are more similar. Consider, for example, “ $2 < 10$ ” and “ $2 < 11$ ”, or “John loves Mary” and “John loves Emma”.

Using the term “truth-determinant” for “*substantial* factor that determines, either alone or together with other factors, the truth-value of a given truth-bearer”,²² let us consider three (distinct yet partly overlapping) degrees of singularity that truth might have:

- (S1) No distinct truth-bearers share *any* truth-determinant,²³
- (S2) Some distinct truth-bearers share *no* truth-determinant,
- (S3) No distinct truth-bearers share *all* their truth-determinants.

Let us also delineate three tasks assigned to a substantive theory of truth.

The task of a substantive theory of truth is to:

- (T1) Informatively account for *the* truth-determinant of *all* truth-bearers,²⁴
- (T2) Informatively account for *some* truth-determinant(s) common to *all* truth-bearers,
- (T3) Informatively account for *some* truth-determinant(s) common to *some* truth-bearers.

Finally, let us call a pair, $\langle S, T \rangle$, a “strategy” for constructing a substantive theory of truth.

The radical disunity argument can be construed as the claim that the strategy required for a substantive theory of truth is $\langle S1, T1 \rangle$ and that this strategy is unviable. Given S1, T1 is unachievable.

Now, if we do assume that the strategy required for a substantive theory of truth is $\langle S1, T1 \rangle$, then the radical disunity argument is right. If the singularity of truth is as strong as S1 says, T1 is not a viable task of a theory of truth. But the assumption that

²² For the present purpose, there is no need for a precise definition of a “truth-determinant”, but an example might be useful. Consider “John loves Mary”. The truth-determinants of this sentence include such things as John’s existence, and/or John’s feelings toward Mary, and/or John’s behavior toward/ concerning Mary, and/or John’s brain states concerning Mary, and/or... i.e., the truth-determinants of “John loves Mary” are the major conditions that have to hold for John to love Mary.

²³ Note: when we use “truth-determinant” in the singular, we allow, though do not require, that it be (finitely) conjunctive.

²⁴ Where “the” is understood in a Russellian manner, i.e., involves an existential requirement. “Of all truth-bearers” implies “common to all truth-bearers”. Note: For the sake of simplicity I am using “the truth-determinant”, in the singular. This formulation can be expanded to a collection of truth-determinants, provided this is done carefully, to avoid trivializing the condition.

$\langle S1, T1 \rangle$ is the right strategy for a substantive theory of truth is utterly unreasonable. First, S1 is wrong about the degree of singularity of truth: sentences like “ $2 < 10$ ” and “ $2 < 11$ ” share *some* of their truth-determinants, as do sentences such as “John loves Mary” and “John loves Emma”. Second, T1 is incompatible not just with S1 but also with the weaker (and more reasonable) S2 and S3. And third, there is a whole array of tasks that can be assigned to a theory of truth which are weaker than T1 yet sufficiently strong (e.g., stronger than T3)²⁵ to make a substantive theory of truth possible. I will call the view that T1 is the right task for a theory of truth the “myth of the common denominator”. T1 says that the task of a substantive theory of truth is to identify *the* common denominator of all truths, but the assumption that there exists some one thing which is *the* substantive common denominator of all truths is unreasonable. And the same holds for the view that there should exist such a common denominator in order for a substantive theory of truth to be possible.

The myth of the common denominator is not unique to truth. Many philosophical subject matters—e.g., knowledge, object, meaning, reference—are associated with concepts for which the idea of a common denominator captured by a single, and often a simple, formula is highly appealing yet also highly problematic and rarely fruitful. (Recall the endless, and eventually fruitless, exchanges of examples and counterexamples over the definition of “knowledge” as “justified true belief”: “*a* knows that *P* iff *a* has a true, justified belief that *P*”.) The problem is that truth, knowledge, object, etc., mark a large and diverse area of action, observation, theorizing, and conceptualization, with partial resemblances and partial unities, but also deep gaps and discontinuities, and therefore an existential assumption of an exhaustive common denominator, i.e., a common denominator that captures the entire substance of truth (knowledge, object), is likely to fail. Today, most epistemologists no longer conceive of their task as providing a definition of (or a necessary and sufficient condition for) the locution “*a* knows that *P*”, but the goal of the theory of truth is still often viewed as a definition of “*P* is true”.

Wittgenstein traces the roots of what we call “the myth of the common denominator” to philosophers’ habit of paying excessive attention to linguistic similarities. Using the metaphor of identically shaped handles of a locomotive, Wittgenstein (1953/58/63, §11–2) warns us against attributing unwarranted significance to the “uniform appearance of words”. All handles in a locomotive, Wittgenstein says, look alike, but they each perform a different function. Likewise (in our case), different sentences “look”, i.e., are syntactically, alike, but they each *perform* a different predication. Just because all sentences of the form “*Rab*” are syntactically alike, their truth conditions do not have to be of the same kind. The truth of a sentence is tied up with its specific function or content (with what properties/relations it attributes to what objects); but the contents of sentences—hence their truth conditions—are too diverse to be systematized by a single principle.

²⁵ And possibly different from T2.

The myth of the common denominator, like the criterial view of truth, “invites” the radical disunity argument and, through it, disillusionment with a substantive theory of truth. But not all views of truth do. The radical disunity challenge arises in the context of some views of truth but not others, and our examination of two of these views, the criterial and common-denominator views, suggests that they are highly unreasonable. In the case of the latter view, we have seen that it is associated with a strategy of truth which is based on unreasonable assumptions. We conjecture that the radical disunity challenge is inherently associated with unreasonable views of truth, and if this conjecture is right, then this challenge is spurious, and deflationism cannot be justified by appealing to it (e.g., by saying that it is needed to avoid this challenge). But if deflationism cannot be justified by appeal to this challenge, then it requires a substantial justification of some other kind, something that has not yet been given.

The unreasonableness of the *radical* disunity challenge does not mean, however, that the nature of a substantive theory of truth is not tied up with *some* disunity challenge. Our next step will be to consider whether there is no other, non-radical yet genuine disunity challenge to theories of truth, a challenge that throws light on the nature of a substantive theory of truth and the direction it can and ought to take.

A MODERATE DISUNITY CHALLENGE

There is a genuine disunity challenge, due to a deep methodological tension between two fundamental principles of theorizing, applicable to many theories, including the theory of truth. These principles are the principle of *generality* or *universality* and the principle of *particularity*, *specificity*, or *differentiation*. They were identified by Kant in the passages on the criterial approach to truth cited earlier, but as we shall see presently, the tension they stand in has nothing to do with the problem of the criterion.²⁶

The principle of generality is succinctly stated by Russell in reference to mathematics: “It is a principle, in all formal reasoning, to generalize to the utmost” (Russell 1919: 196). Yet generality is empty without specificity: “The world”, as Austin puts it, “must exhibit (we must observe) similarities and dissimilarities . . . : if everything were either absolutely indistinguishable from anything else or completely unlike anything else, there would be nothing to say” (Austin 1950: 115). Theorizing requires both generalization and differentiation, yet the two stand in a fundamental tension.

The tension between generality and specificity is especially severe in the case of truth due to the confluence of two circumstances: on the one hand philosophers have set up extraordinarily high standards of generality for our truth theories, culminating in the common-denominator (single-definition, necessary-and-sufficient condition) principle. On the other hand, the concept of truth is extraordinarily broad, complex, and diversified, interwoven in different ways in different areas of our cognitive life, and

²⁶ Later on we will discuss another (non-criterial) version of this tension due to Kant.

applicable to sentences of different kinds: physical, psychological, mathematical, ethical, etc.²⁷ As such it resists any attempt at a simple, sweeping characterization. We may say that the problem of truth is the problem of two forces exerting their pull in opposite directions: the enormous breadth of truth calls for a highly general theory; but the enormous diversity of truths (truth conditions) calls for a multiplicity of narrow and highly specialized theories.²⁸

Now, whether, and how, the two forces can be balanced depends on two things: (i) the degree of diversity of truth (if truth turns out to be diversified in the extreme, each truth will require its own theory), and (ii) our willingness to replace our rigid standards of generality for theories of truth by more flexible and “realistic” standards.

To understand the nature of the moderate disunity challenge let us examine it in a different context, that of science. Many of the issues raised by the ongoing debate on the disunity of science²⁹ apply to substantive theories in general, regardless of how empirical or abstract they are. And the theorist of truth might be able to draw lessons from the resolution of this challenge in science. This approach is especially natural from our general epistemic perspective, given its emphasis on interconnections between all branches of knowledge (holism) and their shared factuality (grounding-in-reality, foothold in the “periphery”).

Challenges to a substantive theory of truth can, in principle, be of three kinds: (a) challenges to substantive theories *in general*, (b) challenges to substantive *philosophical* theories, and (c) challenges to substantive theories of *truth*. By approaching

²⁷ We could classify the different types of sentences in other ways as well, for example, concrete vs. abstract.

²⁸ One objection to this analysis is that it can be applied to the definition of any concept whatsoever. Take, for example, the concept of *accurate portrait*. One could argue that because different people look different, an adequate definition of “accurate portrait” is impossible. But surely such a definition is possible. Here is one: “A portrait is accurate iff it looks like the person it is a portrait of”.

First, I would like to make clear that by “a substantive account of accurate portraiture” I understand something very different from the definition proposed above. A substantive account of accurate portraiture provides the kind of information that a student of portraiture has to assimilate in order to produce accurate portraits, i.e., information about what features of the portrayed object have to be preserved by an accurate portrait (color, shape, proportion—do; size and texture—do not), information about the faithful representation of facial expressions, information concerning the representation of physical surroundings, if any (e.g., the relations between distance, size and proportion), etc.

Second, in spite of the fact that a substantive account of accurate portraiture is far from simple, I do not view the task of constructing such an account as raising the same problems that a substantive account of truth does. This is because essentially the same principles are involved in the accurate portrayal of all persons. To see under what conditions problems similar to those of truth arise, consider the possibility of extending the notion (and practice) of accurate portraiture to objects in general, including objects different from the “mid-sized” physical objects we usually associate with accurate picturing: e.g., heavenly bodies, subatomic particles, sets, numbers, emotions, social institutions. Clearly, the principles underlying the pictorial representation of subatomic particles are very different from those underlying the pictorial representation of human faces. It is at this level of generality—the level of conceiving a *universal* system of pictorial representation—that problems analogous to those we are now considering with respect to truth arise.

²⁹ See e.g., Anderson (1972), Fodor (1974), Suppes (1978), Dupré (1993), Hacking (1996), Cartwright (1999), Laughlin and Pines (2000), and Laughlin et al. (2000).

the moderate disunity challenge to truth through the analogous challenge to science we treat it as falling under the first category.

The disunity of science is alternately conceived as concerning either scientific theories or their subject matter—nature. As in the case of truth, these two types of disunity are closely connected. The disunity of nature challenges science's ability to systematize its subject matter in a general manner due to "the disorder of things" (to borrow the title of Dupré's 1993 book). It is an open question how much order there is in nature and, as a result, whether nature can be subsumed under general laws. The disunity of theories challenges the vision of a unified science. This challenge is often attributed to our cognitive limitations and other factors concerning the mind, but the complexity/disorder of the world plays a significant role as well. Both challenges operate on three levels: total unification (i.e., construction of a "theory of everything"), intertheoretic unification (e.g., reduction of psychology to biology), and intratheoretic unification (e.g., elimination of the particle-wave duality in physics). Among the pertinent considerations involved in these challenges are:³⁰

a. *Complexity of the world:*

- (i) The world exhibits different complexities and interdependencies on different levels, and at each level of complexity entirely new properties appear (Anderson 1972).
- (ii) There are both higher and lower organizing principles, and in the course of investigation we sometimes have to add new levels of basic entities, concepts and principles [Dupré 1993, Laughlin and Pines 2000, Laughlin et al. 2000].
- (iii) The behavior of objects and properties is sensitive to a multiplicity of factors governed by multiple principles (Dupré 1993, Cartwright 1999).

b. *Limitations of our cognitive powers:*

The great complexity of the world on the one hand, and our cognitive limitations on the other, limit our ability to comprehend it by a single, unified principle or theory (Comte 1830).

c. *Partiality of knowledge:*

Human knowledge is, by its nature, partial:

- (i) universal principles and explanations are, due to their high level of abstraction and idealization, inherently partial, covering certain aspects of the phenomena under discussion while leaving others uncovered;
- (ii) as concepts and knowledge are expanded, universal principles become partial, so that what was thought, at one point, to be complete, turns out,

³⁰ The points listed below constitute a pastiche of issues raised by multiple authors. Many of them are related to the issues raised in the beginning of this chapter as well to those raised in the account of our model of knowledge in Part II. At the end of each point I will indicate some of its sources.

later on, to be unfinished (Suppes 1978, Chomsky 1988, Hacking 1996, Hawking 1980, Laughlin and Pines 2000).

d. *Richness of interests, multiplicity of perspectives, human creativity:*³¹

There are many legitimate ways of dividing things into units and many advantageous points of view on things; these, given our creativity, give rise to a multiplicity of theories, a multiplicity which enriches our understanding rather than impedes it (Fodor 1974, Suppes 1978, Dupré 1993).

e. *Other methodological considerations:*

- (i) In any field of knowledge we are continuously confronted with new situations and new problems; as new problems arise, new theories and new methods of investigation are often required (Suppes 1978, Hacking 1996).³²
- (ii) There can be many unifiers, of different interest, in any field of knowledge; hence, there may be room for more than one unifying theory (Hacking 1996).
- (iii) Generality is not always a guide to a better theory. A small collection of simple principles is preferable to a single, highly complex, universal principle (Hacking 1996).

None of these considerations has to do with the highly *empirical* nature of science, and all are either directly applicable or easily extendable to other branches of knowledge, in particular philosophy, due to its broad, diverse, and highly complex subject matter. In the field of *truth*, analogous considerations change our current perspective on existent theories, suggest new developments, and point to new solutions to old problems. Here are a few examples:

Consider the equivalence schema

(E) It is true *that* *p* if and only if *p* (Horwich 1990/8: 7).³³

Deflationists claim that the equivalence schema exhausts the topic of truth. But from the present perspective the equivalence schema describes only one, high-level, principle of truth, and as such it provides a *partial* account of truth, to be complemented by other accounts, centering on other principles of various levels of generality.

³¹ See also discussion in Chapter 4. ³² Partly a citation from Suppes (1978: 14–15).

³³ (E) is closely related to Tarski's T-Schema:

(T) *x* is true iff *p*,

where "*x*" stands for a name of a sentence *S* and "*p*" for a sentence that says the same thing as *S* (see Tarski 1933: 187–8). It is also closely connected to the disquotation schema

(DIS) "*p*" is true iff *p*,

where "*p*" stands for a sentence *S* and "*p*" for a name of *S*.

An example of a somewhat lower-level account of truth is given by Tarski (1933). While the equivalence schema treats all sentences on a par, Tarski's theory distinguishes them along a specific parameter—*logical structure*. Sentences exhibiting distinct logical structures have different Tarskian truth-conditions, whereas those exhibiting the same logical structure have essentially the same truth condition. From the present perspective, Tarski's theory provides an account of a special factor of truth—the *logical factor*—and as such it is a mid-level theory. Most importantly, Tarski's theory, like the equivalence schema, offers a *partial* account of truth (logical structure is not the only thing that determines the truth-value of sentences), and a "complete" account of truth must go beyond it.³⁴

One way of going beyond Tarski's theory was suggested by Field (1972). Whereas Tarski's theory does not distinguish the truth conditions of atomic sentences in any informative manner, Field envisions an informative account of their truth-conditions, based on a physicalistic theory of reference and satisfaction (for the non-logical vocabulary).³⁵ The realization of Field's vision has, however, come upon great difficulties, and these are commonly attributed to a special feature of his approach, namely, its physicalistic orientation. The moderate disunity perspective points to a different explanation. The problem, from this perspective, lies in an implicit assumption of Field's project, independent of physicalism, namely, that the satisfaction conditions of the entire vocabulary of human thought (minus a small part of this vocabulary—the logical vocabulary) are based on one and the same principle, or kind of principle.³⁶ From that perspective the assumption is unwarranted. The totality of extra-logical expressions exhibits an enormous diversity, and one cannot take it for granted that physical and moral expressions, biological and philosophical expressions, psychological and mathematical expressions, expressions pertaining to religion and expressions pertaining to technology, are all governed by the same satisfaction principle, or even by satisfaction principles of the same kind.

Field's reaction to the difficulties facing his project was a retreat to deflationism (though he has been an ambivalent supporter of this approach). But deflationism is not the most productive lesson from these difficulties. A more productive lesson is an *open and flexible methodology*. The theorist of truth, like the scientist, must adjust her methodology to the peculiarities of her subject matter, and to do so effectively she should take pertinent cues from her colleagues in science.

The solution to the moderate disunity challenge facing the theorist of truth is, thus, methodological. The theorist of truth ought to follow straightforward,

³⁴ For a discussion of Tarski's theory see Chapter 8, Section 5.

³⁵ Field uses the term "application" for satisfaction and includes the term "fulfillment" for functional reference, i.e., the value of a functional expression for an argument.

³⁶ With respect to the logical vocabulary, Field believes the task was successfully completed by Tarski (1933).

commonsensical, workable ground rules similar to those followed by the scientist. These include, in the case of truth, such guidelines as:

A. *In constructing a theory of truth, do not legislate in advance the form the theory shall take.* Whether the study of truth will lead to the discovery of a single, universal principle (definition, schema, necessary and sufficient condition), or to a number of partial principles, is an open question, the answer to which largely depends on features of our subject matter, i.e., on things that will emerge in the course of, not prior to, investigation. The answer also depends on our resources and capacities, but they, too, cannot always be determined at the outset of inquiry.

B. *Do not think of the study of truth as focused on a single problem.* Truth is a broad, complex, and diversified topic, and as such it poses a plethora of problems rather than a single problem. Today, it is common to center the study of truth on the subject of truth conditions, but although this undoubtedly is a central subject, it does not exhaust, or even come close to exhausting, the topic of truth. Other subjects include the normativity of truth, the role of truth in knowledge, the relation between truth and correspondence, skepticism and relativism with respect to truth, the interplay between mind and world in creating a standard of truth, and so on.

C. *In developing a theory of truth, aim at a fruitful balance between universality and particularity, similarity and diversity, abstraction and attention to detail, systematicity and applicability.* Dyson's dictum that "every science needs for its healthy growth a creative balance between unifiers and diversifiers" (Dyson 1988: 47) applies not only to scientific but also to philosophical theories.

D. *Think of the development of a theory of truth as a dynamic process, in the course of which the theory is likely to expand, contract, undergo revision, change direction, stall, make leaps of progress, yield unexpected results, etc.* The question is not whether the theory of truth *is* (atemporally or eternally) a unified or a disunified theory; the question is what steps can we take to increase its unity without sacrificing its substantiveness.

These guidelines suggest that we need not conceive of the theory of truth as a Tarskian definition, or a schema, or a criterion, or a common denominator (necessary and sufficient condition). Tarski's definition of truth with its emphasis on logical structure is extremely fruitful in logic, where it is incorporated in the definition of logical consequence and, in this way, makes an invaluable contribution to logical semantics. But it is not clear that to understand whether, why, and how truth applies to, say, ethics, we need a Tarskian definition of truth for moral discourse. A similar point applies to the common correspondence conception of truth. Clearly, if we start with the usual paradigms of correspondence associated with simple observational statements ("Snow is white", "Grass is green", "The cat is on the mat"), we are likely to conclude that correspondence is out of question in

ethics or mathematics. But if we investigate the moral and mathematical domains without prejudice, if we seek to understand them from the perspective of their own philosophies, we open up the possibility of new insights into correspondence, insights that would liberate us from the naive view of true thought as a *mirror of* (or as *isomorphic to*) reality.³⁷

The comparison with science is also relevant to the deflationism-substantivism debate. Consider a scientific theory that limits itself to a few truisms about its subject matter. It says very little about its subject matter, and what it says is neither controversial nor contestable. Such a theory violates some of the most basic standards of a worthwhile scientific theory: informativeness, explanatory power, and substantial justification (refutation). These standards are independent of the empirical character of science, and as such apply to any theory aiming at knowledge. But a deflationist philosophical theory violates these standards. Such a theory has some advantages: it says so little about truth, and its interest in truth is so narrow, that it is compatible with almost anything that can be said about truth. But is this really an advantage? McDowell is known for warning against confusing *justification* with *exculpation*. Saying so little about truth that one cannot be wrong about anything substantive grants one immunity to blame. But what a good theory needs is justification. A theory has to say enough so that on the one hand it is open to critical examination (hence refutation), and on the other hand it teaches us something new about its subject matter—in the present case, the structure and nature of truth, given the complexities of the human cognitive situation. This can only be achieved by a substantive theory. What form a substantive theory should take, however, depends partly on the diversity of its subject matter, and in the case of many philosophical theories this requires some degree of plurality. There are, however, considerable variations in the degree of diversity of different philosophical fields. The subject matters of ethics or logic, for example, are relatively unified, and as such require a lower measure of plurality. But the subject matters of the theories of truth, knowledge, meaning, and ontology are more diversified, and thus require a greater degree of plurality.

The moderate disunity challenge was already recognized by earlier philosophers, in one form or another. To further clarify the present view, let me indicate a few points of similarity and difference with three of those.

(a) *Kant*. In an appendix of the first *Critique*, “The Regulative Employment of the Ideas of Pure Reason”, Kant says:

The logical principle of genera, which postulates identity, is balanced by another principle, namely, that of *species*, which calls for manifoldness and diversity in things, notwithstanding their agreement as coming under the same genus, and which prescribes to the understanding

³⁷ Compare with Chapter 4, Section 1.

that it attend to the diversity no less than to the identity. This principle (of discriminative observation, that is, of the faculty of distinction) sets a limit to possible indiscretion in the former principle (of the faculty of wit [the faculty by which we determine the universal appropriate to the particular, in contrast to the faculty of judgment, by which we determine the particular that accords with the universal]); and reason thus exhibits a twofold, self-conflicting interest, on the one hand interest in the *extent* (universality) in respect of genera, and on the other hand in *content* (determinateness) in respect of the multiplicity of the species (Kant 1781/7: A654/B682).³⁸

Later on he describes this duality as one of “homogeneity” vs. “variety” or “specification”:

Reason thus prepares the field for the understanding: (1) through a principle of the *homogeneity* of the manifold under higher genera; (2) through a principle of the *variety* of the homogeneous under lower species. . . . These we may entitle the principles of *homogeneity* [and] *specification* . . . of forms. (Kant: A657–8/B685–6).

And referring to the principles in (1) and (2) above as “laws”, he elaborates:

The first law . . . keeps us from resting satisfied with an excessive number of different original genera, and bids us pay due regard to homogeneity; the second, in turn, imposes a check upon this tendency towards unity, and insists that before we proceed to apply a universal concept to individuals we distinguish subspecies within it (Kant: A660/B688).

Like Kant, we regard the disunity problem as a problem of balance. But Kant’s construal of the problem is narrower than ours. Kant tries to fit the disunity problem into a neat and orderly picture of our system of knowledge as a hierarchical, species-genera structure, but the tension between unity and diversity, as we see it, is more complex, intricate, and multidimensional than the species-genera picture suggests. Accordingly, we view our system of knowledge as a polymorphic, holistic structure.

(b) *James*. In his lectures on pragmatism (1907), James subsumes the disunity problem under the classical puzzle of “the one and the many”:

If . . . we talk in general of our intellect and its needs, we quickly see that unity is only one of them. . . . What our intellect really aims at is neither variety nor unity taken singly, but *totality*. In this, acquaintance with reality’s diversities is as important as understanding their connexion (James: 59).

The . . . point is to notice that the oneness and the manyness are absolutely co-ordinate here. Neither is primordial or more essential or excellent than the other (James: 62).

While James, too, regards the disunity problem as a problem of balance, his work, as commonly interpreted, suggests a fundamentally unbalanced solution to this problem. James recommends a shift toward radical pragmatism, a pragmatism that requires, at least in philosophy, rejecting the rational, the abstract, and the theoretical

³⁸ Explanation in square brackets based on translator’s fn. 1.

in favor of the experiential, the concrete, and the practical.³⁹ We agree with James that balancing the demands of unity and diversity requires a certain amount of pragmatic “juggling”, but there is no reason that this would conflict with an overall rational, abstract, and theoretical outlook. The disunity challenge is a challenge to the *design* of theories, not to theorizing itself; it necessitates the introduction of some pragmatic considerations into philosophy, not the elimination of theoretical, rational, and abstract considerations from philosophy.⁴⁰

(c) *Wittgenstein*. Wittgenstein’s disunity challenge is expressed in his “family resemblance” remarks. Although Wittgenstein did not use the epithet “myth of the common denominator”, his “family resemblance” remarks can be viewed as a denunciation of this myth. Many concepts, according to Wittgenstein, have no single defining characteristic (or a small set of such characteristics); instead, they exhibit “a complicated network of similarities overlapping and crisscrossing: sometimes overall similarities, sometimes similarities of detail” (Wittgenstein 1953/58/63: §66). This is the case, in particular, with concepts representing the traditional subject matters of philosophy: language, thought, number and, significantly, truth.⁴¹

We share many of Wittgenstein’s views concerning philosophy: that philosophy tends to produce “network” rather than “one characteristic” concepts, that its theories face a serious disunity challenge, that its vulnerability to disunity is due to the breadth, complexity, and multidimensionality of the problems it seeks to resolve, etc. We also support Wittgenstein’s conclusion that philosophers should increase the element of “looking” in their investigations:⁴² “Don’t say: ‘There *must* be something common, or they would not be called [‘truths’]’—but *look and see* whether there is anything common to all” (Wittgenstein: §66).

Wittgenstein, however, goes too far. Like James he concludes that philosophy “may not advance any kind of theory”, that it “must do away with all *explanation*” (Wittgenstein: §109),⁴³ and that it should abandon any aspirations to systematicity. These conclusions, however, are non sequiturs. It is an open question how much disunity there is in various branches of philosophy, whether this disunity rules out the existence of structure (hence of theory and explanation), and how resourceful philosophers can be in facing this disunity.⁴⁴ The tensions between unity and plurality, generality and particularity, abstraction and attention to detail challenge philosophy but need not stifle it. In fact, they create a fertile ground for the

³⁹ A contemporary philosopher who does not view James’s pragmatism as constricting in this way is Ben-Menachem (1995). Be that as it may, the prevalent association of pragmatism with a negative attitude towards abstract, rational theorizing justifies making the point that an adequate solution to the disunity problem does not require such an attitude.

⁴⁰ A similar point was made by Brandom (1994b) with respect to Putnam’s pragmatism.

⁴¹ See e.g., Floyd (1995). ⁴² See Chapter 6, Section 1.

⁴³ See also §126: “Philosophy simply puts everything before us, and neither explains nor deduces anything”.

⁴⁴ This point was also made by Baker and Hacker (1980): 327.

construction of theories, for what is it to construct a theory but to systematically connect elements that, prior to the construction, are disparate, varied, disorderly, and disconnected? Just as importantly, Wittgenstein's contrast between *looking* and *thinking*—"don't think, but look!" (Wittgenstein: §66)—is unwarranted. Much of thinking is looking, and in many fields (for example, metalogic) looking is abstract (e.g., looking at a proof system to see whether it is complete). Philosophers ought to increase the element of "looking" in their methodology: examine the objects of their inquiry, be open-minded and non-dogmatic, aim at correctness, provide justification, be mindful of counterevidence, and so forth. But this does not conflict with theory and explanation.

The lesson to be drawn from the existence of subject matters involving a "complicated network of similarities" and "relationships" (Wittgenstein: §66) is "*Look and think!*". Any knowledge, including philosophical knowledge, requires both looking and thinking. Indeed, looking itself is a complex activity, one that often involves thought as much as, or even more than, sensory perception. Thus, Russell's discovery of a paradox in Frege's logic, Gödel's discovery of the incompleteness of arithmetic, and Tarski's solution to the Liar paradox, involved a considerable amount of *thinking qua looking*, or *looking-thinking*.

The moderate disunity challenge suggests a *moderate pluralism* with respect to the theory of truth. The idea is that the theory of truth may profitably be constructed as a *family* of theories, rather than a single theory. Each theory in the family would investigate some area, aspect, or factor of truth, and together these theories would produce (in the ideal limit) a comprehensive account of truth. This pluralism is moderate since on the one hand it does not rule out the possibility that a single, exhaustive, and substantive theory/principle of truth can, despite the difficulties, be constructed; on the other hand it holds that such a theory/principle (though desirable) is not a *sine qua non* for a thorough and genuine (i.e., substantive) understanding of truth. This methodology should not be confused with a *piecemeal* methodology. The family of theories of truth is not an haphazard array of disconnected theories. On the contrary: the family of theories of *truth*, like the family of theories of *nature*, aims to establish general and systematic connections between seemingly disconnected phenomena; not, however, by oversimplifying complexities and differences, but based on a genuine understanding of common principles and significant interrelations.

Truth, on our conception, is a multifaceted "phenomenon". The *truth* of a sentence, like the *nature* of an event in the world, is the outcome of many factors acting on different levels alongside, and sometimes against, one another. There is no *one* theory of nature (*the one and only* theory of nature). There is a large diversity of theories of nature, each investigating some aspect(s), some force(s), some significant factor(s) in the working of nature. This observation does not invalidate the ideal of *unification* in science. Ever since Thales reduced all multiplicity in nature to one element, the search for a universal principle of nature—a common denominator of

all natural phenomena, or better yet, *the* common denominator—has been one of the main driving forces of science. But it is important to note, first, that scientists are not trying to find an ultimate principle of nature “at all costs”, in particular, not at the cost of *trivializing* science. Second, the viability of natural science is not dependent on the discovery of such a principle. (A “grand” unification principle is a *desideratum*, not a *precondition*, of science.)

I will not attempt to draw a systematic “map” of the family of theories of truth (or the different chapters in the theory of truth) here. In fact, I doubt that such a map can be drawn prior to the actual investigation of truth. But some topics that are likely to appear on such a map are (i) truth conditions of single sentences and of whole theories, (ii) inherently philosophical types of truth: moral truth, metaphysical truth, logical truth, etc., (iii) general philosophical categories of truths: abstract truth, empirical truth, etc., (iv) reducibility of some kinds of truth to others, e.g., of mental truth to physical truth,⁴⁵ (v) relations between truth and other subjects of philosophical inquiry: truth and knowledge, truth and rationality, truth and ontology, (vi) special questions concerning truth: Is the correspondence approach to truth committed to a picture theory of language (the view that language is a *mirror* of reality)? Is the molecular approach to truth (characteristic of all theories associated with Tarski) compatible with a holistic epistemology? What is the relation between the normative and the descriptive element in truth? etc.

And this moderate pluralism with respect to truth does not conflict with its unity. The idea of unity or unification is not restricted to a so-called “ultimate principle”. Partial unificatory principles play a central role in science, and it is here, in the pursuit of greater and greater, yet still substantive, generalities, that lies a key to our solution. This is especially true for one central member of the family of theories of truth, namely the theory of *truth conditions*.

It is this part of the theory of truth for which our earlier terminology of “determinant” (or “substantive factor”) and “strategy” ($\langle S, T \rangle$, where S is the degree of *singularity* of truth and T the *task* of a theory of truth) was designed. Using these notions, we are now in a position to characterize a preferred strategy of a substantive theory of truth conditions (determinants of truths), $\langle S^*, T^* \rangle$. Our analysis suggests the following view of S^* and T^* :

- (S^*) The degree of singularity of truth is intermediate: in principle, not all truths (truth bearers) share all their determinants (significant factors), but there are significant collections of truths with significant shared determinants.
- (T^*) The task of a substantive theory of truth conditions is to account for all the determinants, or types of determinants, of truth in all philosophically significant collections of truth-bearers.

⁴⁵ Field’s 1980 project falls under this category.

It is significant that the present versions of S and T have a larger qualitative component (reflected in the multiple uses of the adjective “significant”) than the versions we have considered in connection with the radical disunity challenge. This is as it should be. Neither the disunity challenge to a substantive theory of truth nor the solution to it is purely (or even dominantly) quantitative. Furthermore, while the adjective “significant” is methodologically important, it is best left unspecified. As a result, the strategy $\langle S^*, T^* \rangle$ is not intended to (and cannot be) used as a precise test for adequate theories of truth conditions. Rather, it provides general methodological guidance that in different cases allows us to assign different weights to different considerations. Yet even in this unspecified form $\langle S^*, T^* \rangle$ has important ramifications. One of these is that the project of a substantive theory of truth is *not* an *either/or* project: the substantivist project does not involve the claim that *either* we construct an absolutely general substantive theory of truth *or else* we abandon the project of a substantive theory of truth altogether. Rather, this strategy allows us to negotiate the generality and informativeness of our theories of truth to the best of our ability, given the conceptual resources available to us on the one hand and the complexities of our subject matter on the other.

II. *Pluralism in perspective*

A number of contemporary philosophers have noted the multiplicity of facets of truth and its impact on the theory of truth: Dummett (1963) regards the realist notion of truth as relative to certain contextual parameters (e.g. discourse about the past, present, and future); Davidson (1969, 1980) correlates various types of truth (truth of sentences with indexicals, truth of sentences about actions and events, etc.) with different extensions of Tarski’s definition; Devitt (1984/91) conceives of the concept of truth as “truth(x)”: “truth(physical)”, “truth(ethical)”, etc.; Resnik (1990) thinks of a variety of bases for the correspondence relation; and so on.

Two philosophers who have made especially important steps toward a comprehensive solution to the disunity problem are Wright (1992, 1999, 2013) and Lynch (1998, 2001, 2004a, 2004b, 2009).⁴⁶ Wright affirms the existence of both universal and non-universal principles of truth, dividing the theory of truth into two parts: a part that deals with its universal principles and a part that deals with its diversified principles. His point is that while the universal principles are central to our understanding of truth, the topic of truth is not exhausted by these principles. These principles are complemented by other, more specific, principles—principles specific to particular domains of truth. Different domains (physics, mathematics, ethics, the comic, etc.) are, or might be, governed by different types of truth: physical discourse, say, by correspondence truth and mathematical discourse by coherence truth.

⁴⁶ See also the authors in Pedersen and C. D. Wright (2013b) and references there.

Lynch is dissatisfied with the *traditional* conception of a substantive theory of truth, and in this respect he is sympathetic with deflationists. But he believes that a *non-traditional* substantive theory is feasible. Unlike deflationists, he holds the view that truth “needs to be substantively explained, not explained away” (Lynch 2004b: 384), and while deflationists’ lesson from the failure of traditional theories is that truth lacks a substantive nature, he says the traditional theories fail to grasp its substantive nature.

Lynch turns Wright’s pluralist theory of truth into a *functionalist* theory, modeled after functionalist theories in the philosophy of mind. This approach enables him to sharpen the pluralist analysis of truth: the concept of truth is a single, high-level concept, defined by its functional role.⁴⁷ This role may be fulfilled by different properties in different domains. Truth supervenes on these properties but is not reducible to them. We may say that truth is differently realized, or has different natures, in different domains. Following Wright, Lynch accounts for the specific realizations of the concept of truth by substantive principles borrowed from competing traditional theories, like correspondence and coherence. Speaking in terms of “realizers of truth”, he suggests (in Lynch 2004a) the following examples of potential types of truth:

- (a) Causal correspondence—might realize truth in the domain of middle-sized everyday physical objects;
- (b) Coherence—might realize truth in the moral and juridical domains;
- (c) Ideal justification—might realize truth in the domain of theoretical science.

While this pluralistic approach to truth represents a considerable advance over more traditional approaches, there are several problems with it. Four of these are: *excessive disunity*, *multiplication of problems*, *difficulty with mixed truths and inferences*, and *shallow unity*.⁴⁸

A. *Excessive Disunity*. The most severe problem with this approach is the extreme disunity it gives rise to. While truth in one field is based on substantive principles of one type—e.g., correspondence—truth in another field is based on substantive principles of an altogether different type—e.g., coherence or ideal justification. Since the differences between correspondence, coherence, and ideal justification as realizers of truth are quite extreme, the result is an inordinately *disunified* theory of truth—a theory in which what it is for, say, a physical statement to be true is completely different from what it is for a mathematical statement to be true.

⁴⁷ In his 2009 book, Lynch no longer characterizes truth as a high-level concept, but he continues to characterize it as a single concept defined by its functional role.

⁴⁸ For a sympathetic yet critical discussion of a variety of aspects of Wright’s and Lynch’s pluralistic theories of truth and another form of pluralism about truth, “Alethic Disjunctivism”, see C. D. Wright (2005, 2010, 2012), Pedersen (2010, 2012a, 2012b), and Pedersen and C. D. Wright (2012/13, 2013a).

B. *Multiplication of Problems.* By selecting correspondence, coherence, ideal justification, and other types of traditional theories of truth as realizers of truth, the pluralist increases the number and severity of problems his theory has to confront. Surely, the pluralist approach eliminates some applicational problems—for example, the problem of differences between fields of truth facing the correspondence theory of truth. But it multiplies the *in-principle* problems concerning the *adequacy* of other conceptions of truth. Thus, the pluralist has to confront *both* the in-principle problem facing the coherentist approach to truth *and* the in-principle problem facing the ideal-justification approach to truth.

Consider, for example, Lynch's argument against the coherence theory of truth. Lynch describes the coherence theory as follows:

According to the coherence theory of truth,...it is coherent believing that makes it so. That is, a proposition is true...if and only if a belief in that proposition would be a member of some coherent system of beliefs. A belief system is coherent, roughly, to the degree that its members are (a) consistent with each other, and (b) display mutual relations of support. On this picture, in other words, beliefs don't fit the facts; they fit with other beliefs (Lynch 2004a: 69).

His main criticism of the coherence theory is an in-principle criticism:

[Our system of beliefs may] all hang...together... But just because it hangs together doesn't mean it is true. It may be true, but its coherence doesn't make it so" (Lynch: 69–70).

His point is that no matter how well a proposition P coheres with a given body of beliefs, the question still arises: "Is P true?". To say that its coherence makes it true, Lynch (rightly) says, is to adopt an excessively weak standard of truth: "coherence theories are too permissive" to capture the idea of truth (Lynch: 69).

But the same objection applies to coherence as a *realizer* of truth. A proposition appearing in a given moral discourse may cohere with all our moral (and other relevant) beliefs, but this by itself does not make it true. Coherence is "too permissive" to serve as a sufficient condition for truth not just in the physical domain but also in the moral domain.

One might reply that coherence satisfies all the universal "platitudes" of truth—those capturing what most people believe about truth prior to philosophical theorizing.⁴⁹ But this would not solve the problem. If the platitudes of truth include what most people believe about truth (prior to substantial philosophical inquiry), then one of the platitudes is: "Truth is not coherence; coherence is too weak to serve as truth". And this platitude, being universal, applies to truth in all domains.

⁴⁹ I will discuss the platitudinous approach to truth later.

C. *Difficulty with Mixed Truths and Inferences.* Consider the sentence:

- (1) This cat is wet and it is funny (Tappolet 2000: 384)⁵⁰

and assume the kind of truth applicable to

- (2) This cat is wet

is correspondence and the kind of truth applicable to

- (3) This cat is funny

is coherence. What kind of truth is applicable to (1)?

Or consider the inference:

- (4) Cruel cats are hungry. This cat is cruel. Therefore this cat is hungry (Tappolet: 383).

Assume the type of truth applicable to the second premise is coherence and the type of truth applicable to the conclusion is correspondence. Assume further that both premises are true. What type of truth is transmitted from premises to conclusion in this logically valid inference? And how can the coherence truth of one sentence play an essential role in guaranteeing the correspondence truth of another sentence?⁵¹

The problem of mixed sentences and inferences is further exacerbated by the fact that correspondence truth-conditions are generally given in terms of reference and satisfaction, i.e., in terms of features of, or relations involving, subsentential linguistic entities, while coherence truth-conditions are generally not given in those terms and do not involve subsentential linguistic entities. How can the two methods be combined in determining the truth conditions of mixed sentences?

Finally, the problem of mixed sentences and inferences also arises for sentences which are not logically complex and inferences which are not logical. Let us focus on sentences and consider two types of mixed sentences: (i) mixed atomic sentences, and (ii) mixed complex sentences whose complexity-inducing operators are not logical. An example of the former might be:

- (5) Causing pain is bad.⁵²

Assume:

- (a) *Causing* is a physical operation, *pain* is a mental entity (state, property), and *bad* is a moral property;

⁵⁰ Here and in (4), my display differs from Tappolet's. Throughout the volume I assign my own numbers to cited sentences.

⁵¹ For proposed solutions and related comments see e.g., Lynch (2004b, 2005), Sher (2005), and Wright (2013). I will get back to this issue in Chapter 8.

⁵² If you prefer to parse (5) as a non-atomic sentence, take any atomic proposition with "mixed" concepts, i.e., concepts from different domains of truth, and apply an analogous line of reasoning to it.

- (b) *Correspondence truth* is applicable to sentences about physical operations, *pragmatic truth* to sentences about mental entities, and *coherence* to sentences about moral properties.

Which kind of truth is applicable to (5)?

An example of a complex, yet not *logically* complex, sentence is:

- (6) Jane's face is beautiful *because* it is symmetric.

Let us assume that (6) is true and that its two sentential components are true as well. If one of these is a coherence truth and the other a correspondence truth, what kind of truth is (6)?⁵³

D. *Shallow Unity*. A more general objection to pluralism about truth is: if truth is based on multiple principles, in what sense are these principles of the same thing, namely, truth? Wright and Lynch meet this objection by introducing universal *platitudes* of truth, satisfied by all its realizers. But there are fundamental problems with establishing the unity of truth on platitudes, one of which is the shallowness of such a unity. We will discuss this problem in connection with the unity challenge below.

III. Unity challenges

UNITY IN SCIENCE AND PHILOSOPHY

The unity challenge, for any theory, is the challenge of finding significant, informative, comprehensive, and enlightening unities in its subject matter. Unity, as many philosophers have pointed out, is a condition as well as a goal of knowledge. One aspect of unity, namely, its contribution to the *explanatory power* of theories, is especially relevant to the present inquiry, since explanation is a central task of substantive theories. The connection between unity and explanation has been emphasized by several philosophers of science (Hempel 1965, 1966, Friedman 1974, Kitcher 1981, 1989, and others), and much of what they say about the role, forms, and problems of unity in science applies to philosophy as well. Philosophical unification, like scientific unification, may be local or global (unification of disparate elements of a single philosophical field, unification of hitherto distinct philosophical fields, unification of all fields of philosophy, and unification of philosophy with other fields of knowledge). It may center on elements of different kinds: laws and principles, theories, arguments, concepts, methods of inquiry, justification procedures, and so forth. Its goals, like those of scientific unification, may vary: generality, economy, explanatory integration of multiple phenomena, etc. Philosophical unification may be stricter or looser: reduction vs. supervenience, supervenience vs. (mere) integration, hierarchical integration vs. holistic integration, etc. Among the reductionist projects in philosophy are idealism, materialism, physicalism, logicism, the *Aufbau* project, the "linguistic turn", and extreme naturalism; all major

⁵³ For further discussion of this and related issues see Chapter 8.

philosophical systems and movements—rationalism, empiricism, transcendental idealism, moderate naturalism, and so on—aim at harmonious integration of some issues, principles, problems, and/or methods.

Two well-known pitfalls of scientific unification are *spurious* unification and *exclusionary* unification. Spurious unification trivially reduces multiple laws to fewer laws without gain in understanding. One example (Hempel and Oppenheim 1948) is conjunction: unifying A and B by constructing their conjunction, A&B. In the field of truth, deflationists often define truth by an infinite list, or an infinite conjunction, of T-sentences (instances of the disquotational schema). Such a definition arguably provides a spurious unification of diverse truth conditions. Field's (1972) criticism of Tarski's theory is also a criticism of (allegedly) spurious unification, directed at Tarski's list-like specification of the truth conditions of atomic sentences.

Exclusionary unification is a flaw in attitude: to think that the success of a scientific theory of "everything" would leave no worthwhile scientific questions unanswered, or would rule out the usefulness of all unifiers outside this theory, is to fall into this trap. Disquotationalism arguably suffers from this flaw as well. It is not uncommon for a disquotationalist to say that *all* there is to truth is disquotation, meaning *all* worthwhile philosophical questions about truth are answerable by the disquotational schema, and the *only* genuine unifier of truth is this schema.

The connection between unity and explanation in science has recently been challenged by Morrison (2000). Morrison argues that the highest degree of unification is achieved by structural explanations, but the best scientific explanations are causal rather than structural. Without getting into the issue of structural vs. causal explanation, her claims that structural explanations are not the only unifying explanations, that different kinds of explanation are unifying in different ways, and that the most general explanations do not automatically impart the greatest degree of understanding, are right. But her implicit suggestion (which does not follow from the above claims) that science admits only one type of effective explanation, namely causal explanation, and her sweeping declaration that "general principles fail to be explanatory in any substantive sense" (Morrison 2000: 33) are unfounded. Her stance against generality, for example, is undermined by the causal principles she presents as examples, since many of these are highly general. The search for a substantive theory is first and foremost a search for explanatory unifiers, on various levels of generality.

SUBSTANTIVE UNIFYING PRINCIPLES

It is often assumed that the universal, unifying principles of truth are thin, obvious, and largely trivial. In contrast, the thick, substantive, non-trivial principles are local—limited to specific types of truth. Wright (1992) and Lynch (2001, 2004b, 2009) characterize the universal principles of truth as "platitudes" or "truisms". They include such statements as:

P1. To assert is to present as true.

P2. Any truth-apt content has a significant negation which is likewise truth-apt.

P3. To be true is to correspond to the facts.⁵⁴

P4. A statement may be justified without being true, and vice versa (Wright 1992: 34).⁵⁵

What is meant by “platitudes” or “platitudinous”? Wright understands “platitudinous” as “superficial” (Wright: 29), “intuitive” (Wright: 34), “uncontentious” (Wright: 34), something “which no one would dispute” (Wright 1999: 226 fn.), etc., contrasting it with “fine-grained” (Wright 1992: 25), “substantive” (Wright: 27), “deep” (Wright: 29), “profound” (Wright: 72), “metaphysically heavyweight” (Wright: 72), etc. Lynch characterizes platitudes in general as commonsensical “folk” principles (Lynch 2001: 732), and in the case of truth, principles belonging to the “*folk theory of truth*” (Lynch 2004b: 392). What principles belong to this theory?

[T]he principles we employ in our folk theory are those the folk tacitly believe, or are rationally committed to. They aren’t those principles that result from technical philosophical argument: thus principles that concern the nature of correspondence, reference, coherence, superassertibility, and the like are not part of our folk theory (Lynch 2004b: 393 fn.).

Elsewhere Lynch characterizes these principles as “truisms” (Lynch 2009: 7) and “preconceptions” (Lynch: 8).

The view that the universal principles of truth are platitudinous rather than substantive, however, is problematic. First: it has *never been established* that truth has no substantive universal principles. Second: limiting one’s interest to the platitudinous aspects of universal principles *weakens the overall substantiveness* of one’s theory. The goal of a substantive theory is to provide a substantive account of its principles, both universal and specific, and this goal is compromised by the platitudinous approach. Third: the universal principles purport to be principles of unity, but the importance of unity lies, to a considerable extent, in its contribution to *explanation* (as we have seen above); hence, unifying the theory of truth by means of platitudes undercuts one of the main points of unification. Fourth: the platitudinous approach is *uncritical*. To make our theories open to criticism we need to endow them with sufficient content and detail, but a theory that limits its principles, or an important subset of its principles, to truisms as such does not meet this requirement. To treat philosophical principles as platitudinous is to presume that there is no possibility of, or need for, a critical approach to these principles, and treating an important set of philosophical principles in this way threatens to undermine philosophy’s claim to being a critical discipline. That is to say, a platitudinous theory, or part of a theory, is frictionless. Ironically, the platitudinous approach stands in considerable tension with Wright’s and Lynch’s otherwise highly critical, substantivist approach.

⁵⁴ As I understand him, Wright (1992, 1999) treats the platitude of correspondence as a non-substantive statement but theories of correspondence for specific fields (e.g., a theory of physical correspondence) as substantive theories.

⁵⁵ P1–P4 are almost verbatim citations.

How, then, shall we reconcile the diversity of truth with its unity? The key to achieving a “fruitful balance” is not platitudinousness but *partiality*. The universal principles are partial in the sense of not exhausting the topic of truth, i.e., leaving room for other principles, on various levels of generality. Recognizing the (*moderate*) *diversity* of truth means recognizing the *partiality* of its commonalities; that is, recognizing that the common principles of truth can provide only *partial* knowledge of truth. But partiality does not imply non-substantiveness: some partial principles are substantive. The tension between unity and diversity in the field of truth does not mean that global substantive principles of truth do not exist. Diversity is perfectly compatible with the existence of substantive global principles, provided they are not required to be exhaustive.

It should be noted that both Wright and Lynch make some steps in this direction. Wright, for example, does not rule out altogether the possibility of substantive unifying principles of truth. Talking about the assertability platitude, he says:

On reflection, it is . . . not necessary to insist that there is *no* suitable notion of deep assertoric content. It suffices that there is, at any rate, at least a more superficial one, carried by surface syntactic features; and that a minimal truth predicate is definable on any surface-assertoric discourse (Wright 1992: 29).

In later work Wright (1999) describes the platitudes as “a priori” and “analytic” rather than non-theoretical, coarsely grained, and unsubstantive; and he acknowledges their controversial nature. Lynch, too, acknowledges the platitudes’ defeasibility and revisability, saying they need to be sorted out and admitted/rejected by the theory of truth itself. Lynch further expands the notion of platitude from what the folk believe prior to philosophical investigation to what the folk “are *rationally committed to believing in virtue of principles they explicitly believe*”, saying that this requires a “significant effort to clarify and state them precisely, and here . . . much substantive work in Sher’s sense remains to be done” (Lynch 2005: 339–40). In still later work, Wright (2013) identifies the platitudes with “axioms” in a “network analysis” (p. 135) of the concept of truth—analysis of the connection between “truth” and related concepts. As far I am concerned, such axioms might just as well be viewed as substantive.

Be that as it may, a substantive theory of truth is concerned with the substantive rather than the platitudinous features of truth, and is committed to providing a substantive rather than a platitudinous account of these features. This means that it aims at universal principles that are theoretical, systematic, explanatory, rich in consequences and applications, arrived at by a thorough and critical investigation, provide a deep understanding of significant aspects of truth, can be given a rigorous and precise formulation, etc. That is, a substantive theory of truth is interested in just what the platitudes leave out: the “theoretical” treatment, the “fine-grained interpretation”, and the “substantive content” of its subject matter. A substantive theory

of truth might turn to platitudes as a temporary starting point—a temporary standpoint on Neurath’s boat or a stage in the pursuit of knowledge in which these principles are held “fixed”. But it strives to “unfix” them and study them with the same care, open mind, and critical attitude as all substantive, theoretical, and fine-grained principles. Is it possible that in the end our investigations will reveal that truth has no substantive universal principles after all? Yes, just as investigation might reveal that everything in nature is accidental and lawless or that the only laws governing nature are trivial and of little interest. In the end, it is an open question how much order, structure, or substantive content there is in any given subject matter. Our hope is that there is enough substantive content and structure to the principles of truth, both special and universal.

Yet what kind of principles will the substantive unifiers of truth be? We have indicated that they will be partial in the sense of not exhausting the principles of truth. But we can say more than that. We envision two types of substantive unifying principles: (A) “*Core*” *unifiers*—unifiers that trace the roots of truth to some general human conditions that “give rise” to truth, and (B) “*Specialized*” *unifiers*—unifiers that focus on a particular factor, feature, or task of truth, one that, due to its indifference to most other factors (features, tasks) of truth, is not sensitive to most differences among truths and therefore can in principle apply to *all* truths. In the next chapter we will identify substantive unifiers of truth of both types: core unifiers—“immanence”, “transcendence”, and “normativity”, and a specialized unifier—“logic-ality”. These are not the only universal principles of truth, but they will partly exemplify our conception of a substantive theory of truth. More generally, our goal is to make first steps toward the development of a theory that satisfies the substantiveness requirement, offers a fruitful balance between attention to unity and attention to diversity, explains the role of truth in knowledge, is incorporated in our dynamic model of knowledge, and can be incorporated in holistic foundational projects.

8

Basic Principles of Truth

Having identified both misconceptions about, and real challenges to, a substantive theory of truth, and having proposed a strategy for meeting the challenges in question, we are ready to make the first steps toward developing a substantive theory of truth, one that both builds upon and further expands our theory of knowledge. We begin with a few basic, universal principles of truth and explore some of their ramifications. These principles differ from those commonly appearing in contemporary works on truth, reflecting a shift in orientation. Our orientation toward truth is *epistemic*—not in the sense in which an epistemic approach to truth is commonly understood, but in the sense explained in our discussion of realism in Chapter 5, Section 4. It is common to view the epistemic approach to truth as identifying truth standards with bona fide epistemic standards: evidence, justification, acceptance at the ideal end of inquiry, etc. In contrast, on our understanding, an epistemic approach to truth is one that focuses on the role of truth in knowledge, and one of its central questions is whether knowledge requires standards that go beyond the narrowly epistemic standards mentioned above. Does our system of knowledge require standards of *correctness*—e.g., of attributing to objects properties they have rather than properties they do not have—in addition to standards of evidence and justification? The conception of knowledge developed in the first two parts of this essay, with its universal grounding-in-reality requirement, its dynamic model of knowledge, its basic realism, and so on—suggests a positive answer to this question. We will begin, however, with a principle that is not specifically associated with our conception of knowledge.

8.1 The Fundamental Principle of Truth

The question leading to the fundamental principle of truth is a semi-Kantian question: under what conditions does truth emerge in our cognitive life, given the basic human cognitive (epistemic) situation?¹ Our basic cognitive situation is that of

¹ (i) In asking this question I do not specify in advance whether I am interested in truth as a predicate, as a property, as a relation, as a standard, etc., or what specific kind of objects play the role of truth bearers. The ontological and semantic status of truth and its bearers is partly dependent on the answer to our basic question, and as far as our present interests are concerned, there is no need to settle these issues in advance.

(ii) In speaking about the basic human situation I am referring back to Chapter 5, Section 1.

creatures with considerable cognitive capacities as well as cognitive limitations, seeking to know the world in its full complexity. The question is what combination of modes of thought available to humans gives rise to a concept or standard of truth in this situation.

1. *Immanence*.² One basic mode of human thought relevant to knowledge is directing our mental gaze at some thing(s) and attributing some property (relation, state, etc.) to it (them). I will call this the *immanent* mode of thought and a thought produced in this mode an *immanent thought*. Such a thought has a subject matter with which it is engaged, and when this subject matter is identified with the world (or some facet of the world, or something in the world) we may say that it is a thought about the world and that to think immanently is to think about the world. Immanence, thus, is both a *mode* and a *property* of thought: a mode of thought when thought is viewed as a mental act, and a property of thought when thought is viewed as a contentful object generated by this mental act.³

The word “immanence” has been used in the philosophical literature in a variety of ways. My use of “immanence” is inspired by one of Quine’s uses of the term. In some of his writings (e.g., 1981b) Quine says that to speak immanently is to speak from *within a theory*, where speaking from within a theory is, typically, saying something about things *outside the theory*, things *in the world* (as distinct from the theory). In my own use, speaking immanently is speaking *in the way* one typically speaks when one speaks from within a theory, namely, speaking about some subject matter, attributing properties/relations to some objects, or saying how the world is. Quine’s use highlights a significant dialectic of “immanence”: “immanent” connotes “being *internal* to a theory”, but “being internal to a theory” signifies “being directed at something *external* to the theory”.

My idea of immanence is also connected to a common conception of *intentionality* or *aboutness*, expressed in such characterizations as:

[Intentionality] is that aspect of mental states or events that consists in their being *of* or *about* things (as pertains to the questions, ‘What are you thinking of?’ and ‘What are you thinking about?’). Intentionality is the *aboutness* or *directedness* of mind (or states of mind) to things, objects, states of affairs, events. So if you are thinking about San Francisco, or about the increased cost of living there, or about your meeting someone there at Union Square—your

² I have used the notion of immanence with respect to knowledge and philosophy in Chapter 4, Sections 1 and 2, respectively. Here I will be more specific about this notion.

³ (i) The idea of thought as a contentful object can be generalized beyond its connection with mental acts, but in discussing the basic cognitive situation it is more natural to focus on the connection between the two.

(ii) I use “thought” as a general term that covers belief, assertion, statement, judgment, cognition, etc., viewed as mental acts, as well as all the above along with propositions, sentences, theories, system of knowledge, etc., viewed as contentful objects. All these are associated with *conceptual* cognition, but some of the things I say about thought apply to non-conceptual cognition as well. The discussion of non-conceptual cognition, however, is beyond the scope of this essay.

mind, your thinking is directed toward San Francisco, or the increased cost of living, or the meeting in Union Square. To think at all is to think of or about something in this sense (Siewert 2002/6: 4).⁴

Immanence, as I understand it, is related to other views as well: Frege's view that "in every judgment... [a] step [is made] from the level of thoughts to the level of the reference (the objective)" (Frege 1892: 64), James's view that "[h]uman thought appears to deal with objects independent of itself" (James 1890: 271), and Wittgenstein's view that "[t]he general form of a proposition is: This is how things stand" (Wittgenstein 1921: 4.5).

My use of "immanence", however, is different from other uses of this term in the philosophical literature, including some of its uses by the authors mentioned above. Thus, in various places Quine (e.g., 1970/86, 1986a, 1995) characterizes immanent statements as restricted to our mother tongue, a given object language, scientific discourse, or naturalistically construed discourse. My own conception of immanence does not impose any of these restrictions. Immanent thought, on my conception, is commonly translanguistic: the principles of general relativity, for example, are immanent in my sense, yet they do not belong to a specific language. Furthermore, immanent thoughts need not be object-language thoughts or lower-level thoughts. Metalinguistic and metatheoretic thoughts are for the most part immanent: they have objects, i.e., lower-level thoughts, and they attribute properties to these objects. Thus, the metalinguistic thoughts "Standard first-order logic is complete" and "The Liar sentence is paradoxical" are both immanent. Similarly, many thoughts outside empirical science are immanent. For example, Kant's transcendental claims about the necessary conditions for the possibility of human cognition are immanent in my sense,⁵ saying something about a specific subject matter: human cognition.

Among the distinctive features of the present use of "immanence" (compared to some of its uses elsewhere) are:

- (a) The objects of immanent thoughts are often external (non-mental) rather than internal (mental), and the thoughts themselves often have (or are intended to have) objective rather than subjective force.
- (b) The category of an immanent sentence cannot be demarcated purely grammatically, i.e., by such linguistic specifications as "declarative sentence". On the one hand, a question-answer pair—e.g., "<Is it raining?>", "<Yes>"—can play the role of an immanent thought (as pointed out by Wittgenstein 1953/58/63: §22); on the other hand, a declarative sentence may turn out to do something other than attribute a property to objects, e.g., express an emotion, as suggested by expressivists with respect to moral statements.

⁴ For recent discussions of aboutness see, e.g., Simchen (2012) and Yablo (2014).

⁵ Though not in Kant's own sense, since Kant limits immanence to experiential concepts, thought, and principles (see e.g., Kant 1781/7: A295–6/B352, A326/B383).

- (c) An immanent statement need not be assertoric. Speaking in Fregean terms we may say that any statement that can be attached to (Frege's) content stroke is immanent. For example, a statement that appears in a non-assertoric position within an immanent statement—e.g., as the antecedent or consequent of a conditional—is generally immanent in our sense.
- (d) The range of immanent thoughts is vast and diverse: immanent thoughts may take the form of a sentence or a theory, they may be about real or fictional objects (events, phenomena), they may belong to any field of knowledge (physics, mathematics, philosophy); they may be contingent or law-like, syntactically simple or syntactically complex (logical compounds, modal compounds, subjunctive conditionals), and so on.
- (e) The attributive nature of an immanent thought extends to logically complex thoughts. In straightforward cases, an immanent thought of the form " $\sim Ba$ " attributes to the individual denoted by " a " the property which complements the one designated⁶ by " B " in the given domain of individuals; an immanent thought of the form " $Ba \& Ca$ " attributes to the individual denoted by " a " the property obtained by intersecting the properties designated by " B " and " C "; an immanent thought of the form " $(\exists x)Bx$ " attributes the property of non-emptiness to the property designated by " B ", and so on.
- (f) An immanent thought may be negative, i.e., say that things are not thus and so. (In view of (e) above, negation can be viewed either as a property-modifier or as a situation-modifier.)
- (g) The content of an immanent statement is not always determined by a literal reading of that statement; in some contexts its content might be determined by a metaphoric, reductionist, or some other indirect reading.
- (h) Immanent thoughts are not limited to statements made in earnest: assertions made by actors on a stage, statements made in the course of flirting, ironic utterances, etc., can be immanent (hence open to examination with respect to their truth).
- (i) The domain of immanent thought is not fixed once and for all; its domain is growing and changing, dynamic rather than static.

The category of immanent thought *determines the domain of potential truth-bearers*, or *truth-apt* thoughts, i.e., the range of thoughts for which truth can serve as a standard or to which the properties of truth and falsehood can, in principle, apply.⁷ This category, as we have just seen, is very broad, and its vastness, diversity, inner

⁶ "Refer", "denote", "designate", etc., are used as synonyms.

⁷ Clarifications: (i) Our category of "immanent thought" is coextensional with *our* category of "truth-apt thought", but *not* with all other uses of the latter category. (For example, it is not coextensional with uses of "truth-apt" which are either coextensional with "assertoric" or "declarative", or imply any of these categories.) (ii) While our category of "immanent thought" is *coextensional* with our category of "truth-apt", it is *not synonymous* with it. This is because we construe "immanence" as having rich content that stands on its own and is not subsumed under "truth-apt".

complexity, and dynamic nature explain, in large part, the breadth, complexity, and diversity of our concept of truth. Indeed, it is these features that are largely responsible for the disunity problem faced by the theory of truth.

Immanence by itself, however, is not sufficient for truth. To focus on the world, to say something about it, is not yet to approach it through the prism of *truth*. To arrive at truth we need to view not only the world but also our thoughts about it.

2. *Transcendence*. A second basic mode of human thought central to knowledge and essential for the emergence of truth is the *transcendent mode*. We transcend a given thought, or domain of thoughts, in order to *reflect* upon it, ask and answer questions about it, set norms or standards for it, challenge it, attribute properties to it, and so on.

Transcendence has fallen into disrepute lately. To say that truth is transcendent, it is claimed, is tantamount to saying that it requires a “God’s eye view” on the world, access to “things in themselves”, etc. “Transcendence”, as I use this term here, has none of these connotations. Transcendence is not something mysterious or super-human. Rather, it is something quite simple and commonplace, like ascending to a *Tarskian metalanguage* (metatheory) or moving “sideways” to *another* (e.g., *background*) *language* or *theory*.⁸ These are the paradigms of transcendence on our approach. Transcending an immanent thought, or a region of immanent thoughts, is casting a reflective look on it from a standpoint external to it yet *within the purview of human thought*. To move to a transcendent standpoint vis-a-vis an immanent thought *t* is to find a perspective from which we can examine *t*, ask questions about *t*, say things about *t*, relate *t* to other things—a perspective which can encompass not just *t* itself but also the objects (the world) *t* is directed at, yet a standpoint in which we are still bound by the usual constraints on human thought.⁹

Transcendence in this sense is so central to human cognition and so important for understanding it, that we ought to reclaim the notion most naturally referring to it, “transcendence”, from those who claimed it as an emblem of what is wrong with various kinds of philosophy—traditional metaphysics (Kant), metaphysical realism (Putnam), etc. To this end, let us introduce two new terms, “HG-Transcendence” and “HH-Transcendence”. Characterizing “transcendence” in general as “going beyond a standpoint X to a standpoint Y which affords a view of X and possibly of other things, like the things that X itself purports to offer a view of”, we may characterize “HG-Transcendence” and “HH-Transcendence” as:

HG-Transcendence: Transcending a *human* standpoint, X, to an ultimate, *Godly*, standpoint Y.

⁸ Note that my use of “sideways” here differs from that in McDowell (1994).

⁹ While here I use Tarski’s hierarchy as an example of transcendence in my sense, in Sher (2016a) I use my conception of transcendence (which I describe without reference to Tarski) to provide a philosophical (as opposed to a formal or technical) basis for Tarski’s hierarchy of languages (theories).

HH-Transcendence: Transcending one *human* standpoint, X, to another, *human*, standpoint Y.

“Transcendence” in Kant’s sense, “God’s eye view”, “view from nowhere”, and “cosmic exile” are all examples of HG-Transcendence. One example of HH-Transcendence is Tarskian transcendence, i.e., ascent from a language L to a meta-language, ML, which characteristically has in view both L itself and those aspect of the world that L talks about, or more precisely, that we can talk about using L. Another example of HH-Transcendence is Kantian transcendentalism—i.e., ascent from a level of thought (cognition, knowledge), T, to a level of thought (cognition, knowledge), TR, in which we talk about the conditions for the possibility of thoughts (cognition, knowledge) at level T. Metaphorically, transcending a thought in an HH-manner is shifting our position within Neurath’s boat, finding a standpoint *on* the boat from which to view, study, attribute properties to, evaluate, improve, or discard the given thought.

Given our holistic, dynamic conception of knowledge and our emphasis on epistemic freedom, and using the intuitive (though not mandatory) locution of “transcendence to a *higher* level of thought”, we may describe HH-Transcendence as satisfying the following conditions:

- (a) There are multiple types and levels of HH-Transcendence open to humans.
- (b) There is no highest level of HH-Transcendence. No matter what level of thought (cognition, knowledge) we stand in, we can always transcend to a higher level of HH-Transcendence.
- (c) The notion of higher-than, as applied to HH-Transcendence, is relative to context: in a context in which we reflect on, or study sociology from a philosophical standpoint, philosophy is higher up than sociology, but in a context in which we reflect on, or study philosophy from a sociological standpoint, sociology is higher up than philosophy.
- (d) Sometimes it is possible to transcend X from within X itself, as in the case of studying the syntax of arithmetic using arithmetic tools (à la Gödel). In this case, one aspect of X (e.g., numbers as representing syntactic units) is higher up than another (numbers as numbers).

Since in this essay we are only interested in HH-Transcendence, our use of “*transcendence*” simpliciter will *always abbreviate* “*HH-Transcendence*”.

Now, it is important to recognize that, although transcendence is a *human* operation, it is a *cognitively powerful* operation. Gödel’s completeness and incompleteness theorems, Tarski’s definitions of “truth” and “logical consequence”, Church’s thesis, Turing’s proof of the unsolvability of the halting problem, the Löwenheim-Skolem theorem, Lindström’s theorems, Kant’s Copernican revolution and his defense of the possibility of genuine scientific knowledge, including causal knowledge and knowledge of laws of nature—are all “transcendent” achievements.

The transcendent standpoint is so pivotal to knowledge that a host of disciplines—philosophy, psychology, sociology, etc.—commonly occupy this standpoint. On the one hand, the very idea of a philosophy, sociology, or psychology of knowledge (language, science, cognition, etc.) requires transcendence; on the other hand, transcendence supplies powerful tools for these disciplines.

The existence of multiple routes of transcendence magnifies our cognitive capacities. Not only can we study one discipline from the point of view of another and vice versa, but we can study each discipline from multiple standpoints. We can study science from a philosophical *and* a sociological *and* a psychological perspective—a threefold increase in our understanding of science.¹⁰

Now, *truth arises in a cognitive environment that includes both the immanent and the transcendent modes of thought*. To arrive at truth we need, first, to have immanent thoughts, thoughts targeting some subject matter—the world in a broad sense, and second, have transcendent thoughts, thoughts capable of evaluating immanent thoughts with respect to their target. It is at this level of thought, a level from which we can observe a given immanent thought in relation to its subject matter, that truth enters into our cognitive life—as a concept, property, standard.

It is important to note, however, that *transcendence does not exclude immanence*. On the contrary, most transcendent thoughts are immanent, as is the case for all the metalogical thoughts of Gödel and others mentioned earlier. Transcendent thoughts attribute properties to their target thoughts, relate them to other thoughts as well as to things other than thought (e.g., their objects), and so on.¹¹

The present conception of philosophy as transcendent, or having a significant transcendent component, is very different from other conceptions of philosophy as transcendent, e.g., the conventionalist conception associated with Carnap. For Carnap (in, for example, his 1950 paper), philosophical and metalogical thoughts are *conventional* rather than *factual*; on our conception, in contrast, they are for the most part *factual* rather than conventional.

Immanence and transcendence by themselves, however, are still not sufficient for truth. By ascending to a higher level of discourse we can do things of many kinds with immanent thoughts: we can ask questions about them, attribute properties and relations to them, set standards (norms) for them, doubt (challenge, ground, justify, refute, enjoy, be disgusted by) them, etc. But while some of these things directly pertain to truth, others do not. We can ask whether a given immanent thought is your thought or mine, say it is a beautiful or an ugly thought, and so on. We can even ask questions about the relation between immanent thoughts and the world that are not questions of truth. For example, we can ask whether a given immanent sentence

¹⁰ Similarly, we can examine a given work of art—say, a novel—both with respect to its literal and with respect to its metaphorical meaning, both with respect to its artistic value and with respect to its political value, and so on.

¹¹ In Chapter 4, Section 2 I talked about the immanent transcendence of philosophy.

names an object it refers to by a word whose sound imitates that object's sound (onomatopoeia), whether it describes a given situation briefly or at length, whether two sentences describing the same situation are synonymous, whether the objects a given theory is about are animate or inanimate, observable or unobservable, etc. *Truth* arises when we ask questions of a special kind about immanent thoughts.

3. *Normativity*. A third basic mode of human thought central to knowledge and essential for the emergence of truth is the normative mode, a mode closely related to humans' proclivity for *critical reflection*. Here is how Korsgaard explains the generation of normative concepts:

Normative concepts exist because human beings have normative problems. And we have normative problems because we are self-conscious rational animals, capable of reflection about what we ought to believe and do. That is why the normative question can be raised in the first place: because even when we are inclined to believe that something is right and to some extent feel ourselves moved to do it we can still *always* ask: but is this really true? and must I really do this? ... It is ... because we are normative animals who can question our experience, that normative concepts exist. ... Normativity is a problem for human beings because ... it is always possible for us to call our beliefs and motives into question (Korsgaard 1996: 46–9).

The normative mode of thought is a mode of questioning, evaluating, accepting, rejecting, setting standards for, recommending, demanding, forbidding, approving an act, behavior, desire, thought in light of what we value, positively or negatively. Normative thoughts are associated with critical questions of the kind indicated by Korsgaard. Here we are interested in critical questions concerning immanent thoughts. There is a large variety of such questions. Three important clusters of critical questions directed at immanent thoughts are:

- (C1) Questions such as: Is it so as a given immanent thought X says it is? Is the world as X holds it to be? Do the objects X refers to have the properties, or stand in the relations, X attributes to them? And so on.
- (C2) Questions such as: Is it justified to assert (believe, hold, accept, include in our system of knowledge) X? Is there adequate evidence for (against) X? And so on.
- (C3) Questions such as: Does X cohere with our body of beliefs (our system of knowledge, the rest of our theories)? Does X stand in a relation of mutual support with the rest of our beliefs (judgments, assertions)? And so on.

I will call the questions in these three clusters the questions of *truth*, *justification*, and *coherence*, respectively. Our interest in this chapter is with the questions of truth or, abstracting from their differences, the question of truth.

The question of truth is the question whether an immanent thought X “measures up” to reality. At issue is whether X is systematically connected to reality in a way that justifies a positive answer to the question of truth as it relates to it. *Truth* is a *standard* for a positive answer to this question, a standard satisfied by some immanent thoughts but not by others. When a given immanent thought satisfies this standard

we say that it is *true*, or that it has the *property* of truth. This duality underlies the *normative-descriptive* nature of truth claims. In as much as such claims are claims about the satisfaction of a standard, they are *normative*, in as much as they are claims about the possession of a property, they are *descriptive*. This duality inheres in all properties and standards. For each standard there is a property corresponding to it (the property of satisfying the standard), for each property there is a corresponding standard (the standard of attributing the property only to objects that possess it).¹² A related duality is also inherent in the nature of *immanent* thoughts. Immanent thoughts are configurations of two components: a representational component which is descriptive, and an attributive component which is normative. These two components are clearly discernable in *quantificational* thoughts. Consider a thought of the form “ $(\exists x)(Ax \ \& \ Bx)$ ”. Its *inner* part, “ $Ax \ \& \ Bx$ ”, delineates, represents, or *describes* a situation with one open parameter (marked by “ x ”), namely the situation of being in the intersection of A and B.¹³ Its *outer* part, “ $\exists x$ ”, performs an *attributive* act, namely that of saying that the situation described by the inner part is not empty, i.e., that there is at least one individual (in the underlying universe) that stands in the intersection of A and B or that has the property of standing in the intersection of A and B.¹⁴ This is another root of the *normative-descriptive* duality of truth: The norm of truth, from this perspective, is a norm of *correct attribution*, the norm of *attributing* to objects properties they *have* rather than properties they do not have (attributing to reality features it has rather than features it does not have).

The normative element in truth can also be characterized in terms of “endorsement”. Thus, speaking about propositions (which fall under our notion of immanent thought), Wright says:

[In one of] its most basic use[s, ‘true’] is essentially a device of endorsement . . . [W]hat is it to *endorse* a proposition? Endorsement generally involves an element of recommendation, or approval of an item as meeting a certain standard. . . . In this way, affirmations of truth—and likewise denials of truth—are normative claims. To endorse a proposition as true is to affirm that it is acceptable as a belief or statement; to deny that a proposition is true is to affirm that it’s correspondingly unacceptable. . . . [I]f ‘true’ is . . . a device of endorsement, then in using it, I’m saying that a proposition is in good shape as far as certain relevant norms are concerned (Wright 1999: 211).

Although truth-claims are both normative and descriptive, philosophically we give the normative view of truth a certain priority over the descriptive view. And while truth, as we have noted above, is not the only norm or standard for immanent

¹² This is partly a Fregean view. See discussion of the normativity of logic in Chapter 10, Sections 5–6.

¹³ Or more precisely, being in the intersection of the extensions of “ Ax ” and “ Bx ” in an underlying universe.

¹⁴ In the case of “ $Ax \ \& \ Bx$ ” the descriptive/representational part is captured by “ $Ax \ \& \ Bx$ ” (as above). The attributive part is captured by “ c ”. We may see this more clearly if we rewrite the sentence as “ $[c]:(Ax \ \& \ Bx)$ ”. The inner part describes a complex property, and the outer part attributes this property to c (the object denoted by “ c ”).

thoughts,¹⁵ truth is one of the most fundamental standards for such thoughts. A theory of truth explains this standard, specifies its principles, and works out its connections to other standards of human thought.

The immanence, transcendence, and normativity principles can be combined into a single core principle of truth. Focusing on some of the most basic elements of truth as it does, I will call it “the fundamental principle” of truth.

4. *The Fundamental Principle of Truth.* Roughly, the fundamental principle of truth says that truth is a normative transcendent standard for immanent thoughts, the standard of measuring up to reality or to the way things are. Briefly, and without purporting to give a complete, precise definition of this principle, we may formulate its content as follows:

(FPT) *It is a fundamental principle of truth that:*

1. *Truth arises in the cognitive environment of three basic modes of human thought: the immanent, transcendent, and normative modes.*
2. *Truth is a transcendent standard or norm for immanent thoughts, a standard of measuring up to reality for such thoughts, or of giving a positive answer to “the question of truth”—the question of whether things are as a given immanent thought says they are—as it applies to such thoughts.*¹⁶

Speaking in terms of “the question of truth”, we may further characterize the fundamental principle by saying that it offers a threefold characterization of each of its constitutive subprinciples:

Immanence:

- (i) Truth is a standard for *immanent* thoughts.
- (ii) The question of truth arises for *all* immanent thoughts.
- (iii) Truth statements—i.e., statements of the form “X is true/false” (“It is true/false that X”, etc.)—are *immanent*.

Transcendence:

- (i) Truth is a *transcendent* standard for thoughts.
- (ii) The question of truth is a *transcendent* question.
- (iii) Truth statements are *transcendent*.

Normativity:

- (i) Truth is *normative* for thoughts.
- (ii) The question of truth is a *normative* question.
- (iii) Truth statements are *normative*.

¹⁵ Coherence, justification, empirical verification, explanatory value, utility, etc. are other standards for immanent thoughts.

¹⁶ (1) applies to justification as well, but the normative question of justification differs from that of truth. I will discuss some aspects of the relation between truth and justification shortly.

The *fundamental principle of truth* identifies some of the most basic elements of truth. But truth itself is a fundamental element of human life, and to understand truth we need to understand this aspect of it as well.

Some of the reasons that truth is fundamental for humans belong to the practical domain. According to Field (1972) and Williams (2002), for example, truth is fundamental for humans because of the practical need for, and advantages of, saying correct as opposed to incorrect things and being truthful as opposed to lying. Other reasons are more abstract. For example, Korsgaard's discussion of normativity (see citation earlier in this section) suggests that truth is fundamental for us because we are creatures who ask, and care about, critical questions in many areas of our life, and the question of truth is a critical question pertinent to most of these areas. In this essay we are interested primarily in epistemic reasons for the importance of truth, and these reasons have to do with the basic epistemic situation. Truth is fundamental for humans because of the confluence of four circumstances: (a) our considerable epistemic ambitions, (b) the considerable complexity of the world, (c) our cognitive limitations, and (d) the intricacy of our cognitive resources.

(a): Our desire to know the world is deep-seated. Since Kant, it is problematic to say that we want to know the world "as it is". But most of us do want to know the world as it is, and we do not identify this with knowing the "thing in itself" or "noumenon" in Kant's sense. Another way to express this is by saying that we are *not* seeking a comforting, or an interesting, or an aesthetically pleasing, or a convenient, or a practically useful story about the world; we are seeking *true* theories about it. What we want are theories that tell us what objects there are in the world, what properties they have, what relations they stand in, what laws they (and their properties and relations) are governed by, and so on. In short, what we want is a *true* theory of the world.

(b): Given the complexity of the world relative to our cognitive capacities, we cannot take it for granted that what we naturally believe about the world, or even what we come to believe following a careful investigation, is not erroneous. We need, therefore, a standard of correctness for our beliefs. We need to transcend our beliefs and theories and measure them against their target, reality, i.e., apply a standard of *truth* to them—a human standard of truth, yet a standard of truth all the same.

(c): The possibility of error, hence the need for a standard of truth (a standard for identifying error), is especially pressing due to our cognitive limitations. The situation is analogous in some (though not all) respects to the situation we find ourselves in concerning *justice*. In *Treatise of Human Nature* (1739–40) Hume raises the question: Why does the problem of justice arise for humans in the first place? His answer is that the problem of justice arises for humans due to certain "deficiencies" in human character on the one hand and the environment on the other, or due to certain "gaps" between them:

[J]ustice takes its rise...as a remedy to some inconveniencies, which proceed from the concurrence of certain *qualities* of the human mind with the *situation* of external objects.

The qualities of the mind are *selfishness* and *limited generosity*: And the situation of external objects is their *easy change*, join'd to their *scarcity* in comparison with the wants and desires of men. . . . [I]f every man had a tender regard for another, or if nature supply'd abundantly all our wants and desires, . . . the jealousy of interest, which justice supposes, cou'd no longer have place; nor wou'd there be any occasion for those distinctions and limits of property and possession, which at present are in use among mankind. Encrease to a sufficient degree the benevolence of men, or the bounty of nature, and you render justice useless. . . . '[T]is only from the selfishness and confin'd generosity of man, along with the scanty provision nature has made for his wants, that justice derives its origin' (Hume 1739–40: 317–18).

A similar answer holds for the question “Why does the problem of *truth* arise for humans in the first place?”. It is because of the gap between our cognitive ambitions and our ability to satisfy them that the question of truth arises for us and a norm of truth is both indispensable and advantageous. Had life been so hospitable that we needed no information about the world, had we no desire to have an intellectual understanding of the world, were we incapable of cognitively diverging from or going beyond what is actually the case in the world, had we no limitations that blocked or distorted information, then a norm, concept, and theory of truth would have been of little use for us. But since our immanent thoughts are bound to diverge—sometimes and/or in some respects—from what is the case about their subject matter, the question of truth, the question whether our thoughts measure up to reality, always arises for us. And since we badly need and strongly want information about the world, this question is vital for us.

(d): Our cognitive capacities, too, make a standard of truth pivotal for us. On the one hand, it is our cognitive capacities which make the pursuit of knowledge possible, and hence a standard of truth useful for us. (Without capacities to acquire knowledge, there would be no point in such a standard.) On the other hand, our cognitive capacities paradoxically contribute to the likelihood of error. Consider imagination. Imagination is an important tool of knowledge, a tool that can help us figure out new ways of reaching those regions of reality which are not easily accessible to us. But imagination can also lead us astray. Imagination makes it difficult for us to distinguish what is real and what is not. A similar point was noted by Freud in connection with “passages of ideas”. Explaining our need for *reality checks*, Freud says that while “in animals” consciousness is always triggered from the outside, “in men . . . internal processes”, such as “passages of ideas”, may also trigger consciousness. And they can do so in a way that is not always transparent, so that humans cannot always distinguish events triggered by external stimuli and those triggered by internal stimuli. As a result, in humans “a special device is called for in order to distinguish between the two possibilities—a device known as *reality-testing*. The equation ‘perception = reality (external world)’ no longer holds” (Freud 1940: 19). The combination of rich and diverse cognitive activities (which make the development of a rich body of human knowledge possible), opportunities for error generated by these

activities, and the ability to correct error using activities of the same kind, makes a standard of truth paramount for humans.

These reasons are quite different from others, often given by deflationists, which focus on certain technical uses of truth-expressions: using the truth predicate as a tool for endorsing sentences one cannot spell out (e.g., “The Löwenheim-Skolem theorem is true”), using truth-locutions as a proxy for substitutional quantification, etc.¹⁷ These uses do not conflict with the reasons given by our account, but they are of a different order of significance.

The fundamental principle of truth is a principle of both freedom and friction. The cognitive setting in which truth emerges has elements of both types: both acts of freedom (turning our cognitive gaze at the world, saying things about it, transcending what we say, asking critical questions) and a constraining world. Furthermore, the truth standard itself, like most standards or norms, is on the one hand a standard freely generated by us, and on the other, a constraining standard.

One might ask in what sense the fundamental principle of truth is a substantivist principle. Our answer is that, first, a successful formulation of this principle will tell us something substantive (interesting, informative, explanatory) about an important subject matter (the roots of truth in our basic cognitive situation). And second, this principle calls for a number of substantive investigations: investigations of the dual contribution of mind and world to the truth conditions of sentences (theories), of the potentially diverse network of “cognitive routes” from mind to reality involved in truth, of the differences and similarities between the ways the standard of truth works in different fields, of the applicability of truth to fields like ethics, of the vulnerability of truth to skepticism, and so on. Of course, by limiting our attention to trivial aspects of the fundamental principle of truth we render it a deflationist principle, but this holds for any substantivist principle.

Some philosophers believe that there is no need for a special standard of truth, distinct from our standard of justification. It is important to understand why this belief is wrong. The question of justification is an important question, and it is related to the question of truth; but it is different from the latter and is not a substitute for it. Indeed, justification itself requires a standard (like truth) to guide it. Justification is relative to goal, interest, standard. Given a set of sentences (a story, a theory, a report, a verdict, etc.), we may justify it with respect to different standards: standards of practical value, beauty, economy, predictive power, explanatory power, simplicity, compliance with our legal code, etc. In the context of knowledge—as it is understood in this essay—the paramount standard is truth. The role of truth in justification is twofold: it tells us what we have to justify—the claim of a given set of sentences to describe the world correctly, rather than its claim to, say, beauty. And it tells us what counts as evidence or justification—counting the number of rhymes in a canonical

¹⁷ See e.g., Horwich (1990/8: 2–5).

formulation of a given theory does not, but checking whether the objects it is about have the properties it attributes to them does. (Similarly, it tells us what counts as a counterevidence or refutation—finding a counterexample to one of its laws does, but showing that testing its claims would be too expensive for us does not.)

By focusing on truth as a standard for knowledge, we put aside some interests that are common in the contemporary literature on truth, for example, our interest in the use of “true” and “truth” in everyday discourse or natural language. Of course, there are significant connections between the role of the standard of truth in knowledge and the use of the word “true” in everyday language, but studying the one is not sufficient for knowing the other. “True” in “‘Snow is white’ is true” might be *redundant*, but that would not make subjecting our theories of snow to a standard of truth redundant.

Although the “fundamental principle” is not the only principle of truth, it has many ramifications. As an example, let us turn to skepticism.

8.2 Ramifications for Skepticism

The skeptical challenge to truth¹⁸ is formulated by Dworkin as follows:

Is there any objective truth? Or must we finally accept that at bottom, in the end, philosophically speaking, there is no “real” or “objective” or... “fact of the matter”... truth about anything, that... [there] are just... convictions, just conventions, just ideology, just badges of power, just the rules of the language games we choose to play, just... instinct[s], imagination and culture? (Dworkin 1996: 87).

As applied to our own theory, this challenge is expressed by questions such as: Is there is an objective standard, or a family of objective standards, of truth for immanent thoughts? Can a significant collection of thoughts (types of thought) in principle satisfy this standard?

The skeptical challenge is naturally divided into two kinds: *local* and *global*. Local skepticism questions the existence, and satisfiability, of a standard of truth in a specific area; global skepticism questions the existence, and satisfiability, of a standard of truth in any area.

The fundamental principle of truth points to an inherent connection between the emergence of truth and the emergence of skepticism. Skepticism with respect to truth, both global and local, is sanctioned by the same principles that sanction truth itself. We can view skeptical challenges to truth as challenges to “truth thoughts”,

¹⁸ (i) The skeptical challenge to truth is part of the skeptical challenge to knowledge. Here I will focus on those aspects of the challenge that center on the former, although the connection to knowledge will always be implicit in the background.

(ii) The philosophical literature on skepticism in general is prodigious. The present discussion and references will be limited to a few pertinent issues and a small number of representative works.

i.e., immanent thoughts focusing on truth. Truth thoughts include thoughts of the form “X is true”, “Some thought is true”, “There is a standard of truth for field X”, “There is a standard of truth for some/every field”, etc. According to the fundamental principle, once an immanent thought arises at level ℓ , critical thoughts, including truth claims, with respect to it arise at level $\ell+1$. And similarly, once a truth-claim arises at level ℓ , critical thoughts, including skeptical challenges, with respect to it arise at level $\ell+1$. These challenge the theorist of truth to refute the claim that there are no truths, or no standards of truth, either in field Y or in any field. In a sense, the question of truth itself has a skeptical element: “Does X (any thought in field Y, any thought whatsoever) *really* measure up to reality?”. And close relatives of this question have an even stronger element of skepticism: “*Can* X (any thought in field Y, any thought whatsoever) measure up to reality?”. It follows that the same setting that gives rise to truth gives rise to skepticism about truth, and that for the same reason that the question of truth itself is central to humans, so is its skeptical counterpart: both are part and parcel of a critical attitude to knowledge. We thus agree with Stroud when he says:

Philosophical scepticism is a ‘benefactor of human reason’ (Stroud 1984: 256).

We also share Stroud’s belief that:

[T]he study of scepticism . . . reveal[s] something deep or important about human knowledge[, human nature[, and] the urge to understand them philosophically (Stroud: ix).

The success of skepticism—both global skepticism and skepticism with respect to particular fields—is an open question, requiring a substantive answer. To further examine skepticism with respect to truth, let us begin with local truth-skepticism and then proceed to global truth-skepticism.

I. Local truth-skepticism: ethics

One example of local skepticism with respect to truth is skepticism with respect to *moral* truth. Based on contemporary literature (e.g., Gibbard 1990) we may view moral skepticism as coming in different *types* or *degrees*. Three types of moral skepticism are:

- (M1) *Moral statements have a standard of truth, but this standard is a standard of a non-moral truth, namely psychological truth.*

This type of moral skepticism is reductive. The idea is that although the surface structure of moral statements is “X has moral property Y”, their deep structure is “Speaker Z has attitude Y^* toward X”, where “ Y^* ” stands for some psychological attitude, mental state, or emotion (e.g., “Y” stands for “good” and “ Y^* ” for “approbation”). The standard by which the truth of moral statements is measured is, therefore, a standard of *psychological* truth.

- (M2) *Moral statements have a standard of truth (either a standard of sui generis moral truth or a standard of some other kind of truth), but this standard is not their primary standard of success.*

This type of moral skepticism is that of irrelevance (or minor significance). Here the idea is that while moral statements do attribute properties of one kind or another to objects (actions, intentions, etc.) and these attributions are subject to a(n appropriate) standard of truth, the main goal of moral statements is to do something else, e.g., express, support, or arouse feelings or attitudes. The primary success standard of moral statements is therefore a standard of persuasiveness or expressivity, not a standard of truth.

- (M3) *Moral statements have no standard of truth whatsoever. They do not attribute any properties to any objects; indeed, they do not say anything at all. All they do is use words to express attitudes or emotions.*

This is a strong skepticism regarding moral truth. It says that moral statements do not attribute any properties to any objects or purport to say how things are, they *merely* express attitudes or feelings, and expressions of feelings cannot be subject to a standard of truth at all. In a sense, this skepticism says that moral statements are not *immanent*. They are not candidates for truth-bearing.

All three types of moral skepticism are compatible with the fundamental principle, according to which only immanent statements are bearers of truth, and even immanent statements are merely *potential* bearers of truth. From the point of view of this principle, it is an open question whether moral statements are literal, whether truth is their main standard of success, whether they are immanent, and if they are, whether they are actual (and not just potential) bearers of truth. Our model of knowledge shows that a standard of truth, a standard of measuring up to reality, is *structurally* possible for highly abstract disciplines (such as ethics, logic, and mathematics, if they are indeed disciplines), since all disciplines, including abstract disciplines, are capable of moving to the periphery. But whether ethics is a discipline (in the sense of a field of theoretical knowledge), what kind of discipline it is, whether it is primarily a discipline, and whether it actually has a standard of truth—all these are open questions, the subject of specific investigations, not just into the nature of truth, but also, and importantly, into the nature of morality.¹⁹ In Section 8.4 I will show how the possibility of a standard of truth for an abstract discipline can be realized in the case of mathematics. In Chapter 10, I will provide an outline that shows, among other things, how this possibility is realized in the case of logic.

¹⁹ An example of a discipline that some might regard (for one reason or another) as having no standard of truth is theology, viewed as a discipline investigating the existence of God (and His properties, if He exists).

II. *Global truth-skepticism*

Global skepticism says that there is no satisfiable standard of truth for any thought whatsoever, i.e., there are no conditions under which the question of truth, as it applies to any immanent thought, can receive a positive answer.²⁰ We have seen how the fundamental principle of truth renders this challenge unavoidable. But this principle also gives rise to (supports, gives precise content to) challenges to global skepticism itself, sharpening our understanding of both the force and the limits of such skepticism.

One way in which the fundamental principle limits the force of skepticism with respect to truth is by supporting Nagel's (1997) claim that global skepticism cannot have the "*last word*". If global skepticism has the last word in the debate, then the last word must belong to some level of discourse ℓ . But if so, the question of truth with respect to the last word (the "final" skeptical claim) arises at level $\ell + 1$, and with it the challenge to establish the last word (the skeptical claim of level ℓ). Structurally, global skepticism, just like truth, cannot have the last word.

The skeptic might deny the satisfiability of the conditions set by the fundamental principle for the emergence of truth. She might make one or more of the following claims:

- (a) There cannot be (or we cannot have) any immanent thoughts.
- (b) There cannot be a transcendent standpoint from which to view our thoughts (any of our thoughts) and say things about them.
- (c) There cannot be a critical standpoint from which to ask (and answer) the question of truth with respect to any thought.

But the truth theorist will then challenge her to justify her claims. And all these claims are extremely difficult to justify. First, these are all negative existential claims, indeed negative *modal* existential claims—claims to the effect that X *could not be* the case—claims that, especially in philosophy (outside logic), are notoriously difficult to establish. Second, in claiming (a), (b), or (c) the skeptic has to confront well-established counterexamples, e.g., the immanence of such statements as "Snow is white", the demonstration—by Tarski—that transcendence (to a metalanguage) is possible, and so on.

These difficulties are further exacerbated by the implausibility of denying that the skeptical claim itself is an immanent thought, that we can—since we did—transcend it and ask whether it is true or false, that there are conditions under which it is false (hence not lacking a truth value). And these difficulties are still further exacerbated by the fact that to establish the skeptical claim the skeptic cannot use any immanent thoughts, any transcendent thoughts, any true or false

²⁰ There might be a more radical global skepticism as well, a skepticism that questions the possibility of an unsatisfiable standard of truth as well, but I will leave this possibility aside here.

thoughts, and so on—something that, to all appearances, makes the task of establishing the skeptical claim self-defeating. If the skeptic says that the skeptical claim is a consistency claim or a *reductio ad absurdum* claim, then the question of truth arises with respect to these claims. And if she refuses to accept our challenge that she justify her claim, why should we accept her skeptical claim as correct, or seriously consider her challenge?²¹

III. Relativism

Another form of skepticism with respect to truth is *relativism*. Different philosophers characterize truth relativism in different ways, and not all versions of truth relativism are skeptical. One version of skeptical relativism with respect to truth says that “standards of . . . truth are . . . relative to language, culture, or biological makeup” or that “tokens of sentences, beliefs or the like are only true relative to a framework” (Swoyer 2003: 1, 24). Philosophers have raised numerous objections to this version of relativism, including the objection that relativism is self-refuting, since it relativizes itself. Another objection is that frameworks (language, concepts, culture, brain structures, etc.) are *conduits* rather than *veils*: they are instruments for *accessing* the world rather than for *hiding* it.²² Both objections resonate with our approach, but since the fundamental principle of truth is more relevant to the latter, let us focus on that objection here. A conduit has its own features, some of which might create certain difficulties for the project it is used for. But this does not mean that these features are bound to undermine the project. We have plenty of resources for overcoming (or mitigating) difficulties caused by the current conduits of truth, and an especially useful tool, suggested by our principle, is *transcendence*. In the first place, transcendence enables us to recognize many features of our language or conceptual scheme; in the second place, it provides us with a standpoint from which we can either eliminate or neutralize some of the difficulties created by these features. In other words, the fundamental principle of truth teaches us that *relativity is in principle transcendable*. We can transcend it by *ascending* to an appropriate metalanguage (or by *moving* to another area on Neurath’s boat) where, by seeing it and “making it explicit” (Brandom 1994a) we neutralize it, free ourselves from it, or control it, at least to a degree. We may call this process “derelativization”. This move enables us to detect both biological and cultural forms of relativity, identify relativity to conceptual schemes, take into account and critically examine relativity to ourselves and to others, and so on. Another thing that transcendence might enable us to do is

²¹ The skeptic might be better off if she limited her challenge to knowledge (instead of truth), and in particular to *absolutely certain* knowledge. But while the foundationalist conception of knowledge is vulnerable to such a challenge, our (foundational-) holistic conception is not. It is not part of our conception that *absolutely certain*, indefeasible knowledge of any kind is needed in order to launch a potentially successful epistemic project.

²² For further objections see Swoyer (2003, Section 5).

introduce improvements into our current system of conduits or develop new, better conduits—either for general or for specific purposes.²³

There is, of course, no final or absolute derelativization of truth, no final or absolute transcendence, no final meta or background language/theory. The expectation that there is is a foundationalist expectation, one that our holistic outlook has freed us from. The important point is that there is always, in principle, a *route* of derelativization. No matter where we stand, it is always open to us to transcend our standpoint, freeing ourselves from its conceptual framework: “S is true relative to standpoint *c*”, made from a standpoint *c**, external to *c*, is not relative to *c*. Of course, what we utter at *c** is also relative to some “coordinates”, *c**’s coordinates, but these, too, can be derelativized by transcending to a *c***, and so on. Relativity claims, like the more straightforward skeptical claims we have considered earlier, *never have the “last word”*. And by exercising transcendence and other forms of epistemic freedom, we can eliminate, neutralize, or overcome cases of relativism that fall under skepticism.

The same strategy is effective in dealing with other problems involving relativism, e.g., *incommensurability* in science. We have already seen how this problem can be dealt with in our outline of the dynamic model of knowledge delineated in Chapter 4.²⁴ There we emphasized the dynamics of fixing and unfixing certain elements of our theories, but the process of fixing and unfixing involves *transcendence*. Given a revolutionary move from one scientific theory to another, the road is always open for transcending both to a level of inquiry from which we may compare and evaluate the two theories with respect to their subject matter, relation to the world, justification, pragmatic desiderata, etc. Not only are we not trapped by a single framework, but we are free to seek better and better standpoints for conducting the comparison.

We have mentioned that some forms of truth relativism do not involve skepticism. One of these is *speaker and assessor relativity*. It says that “[t]o be a relativist about truth is to allow that a sentence or proposition might be assessment-sensitive: that is, its truth value might vary with the *context of assessment* as well as the context of use” (MacFarlane 2005: 321). From our perspective, this is an additional aspect of a substantive account of truth. A substantive understanding, and evaluation, of truth claims requires taking a variety of parameters into account, including the points of view of both speaker and assessor. And what makes it possible to take these points of view into account is *transcendence*. It is important to note that by taking the speaker’s and assessor’s points of view into account we do not weaken the objectivity of our truth claims. On the contrary. Consider aesthetic claims. To say that the tree looks beautiful when viewed from the house and ugly when viewed from the swimming pool is not to say that its aesthetic status is not an objective matter. To say that drinking full-fat milk is healthy under some conditions but not under others is not to

²³ The deliberate development of a new conduit for a special cognitive purpose was exemplified by e.g., Frege (1879).

²⁴ See Chapter 4, Section 2.

say that the nutritional value of full-fat milk is not an objective matter, or that the truth value of statements about health is not objective. In general the operation of transcendence does not weaken the objectivity of our truth discourse; on the contrary, it increases it.

Back to skepticism, let me conclude with two general comments:

(A) Although every truth claim is subject to skeptical challenges and every skeptical claim is subject to truth challenges, the result is not necessarily a draw. At each stage in the debate, one side can have the upper hand either by mounting a compelling criticism of the other side's arguments or by making a positive case for its own claim. In this duel, however, the truth advocate has one important advantage over the skeptic: the "constructive advantage". The truth advocate can win the battle constructively, by developing improved methods of discovery, evidence, and justification, and using these to construct a significant body of theory satisfying pertinent friction requirements, including a strong truth standard, one that requires a significant connection between theory and world. This would demonstrate, in a positive manner, the satisfiability of a non-trivial standard of truth. It is, of course, open to the skeptic to use the unavoidable setbacks in the development of such a body of theory to strengthen her own claim. But as in the case of existence and possibility, where existence is sufficient to establish possibility while non-existence is not sufficient to establish non-possibility, so in the present case, success in the project of truth (knowledge) is sufficient to establish the possibility of success, while setbacks and failures are not sufficient to establish its impossibility.

(B) Skepticism with respect to truth (knowledge) is often motivated by what we may call the "illusory requirement of the absolute". This illusory requirement, which we have already encountered (if not by name) in our discussion of foundationalism and realism (Chapter 2, Section 1 and Chapter 5, Section 4), enters into skepticism in several forms. First, there is the belief that nothing short of an absolute ban on circularity and infinite regress is required for establishing the possibility of truth. Second, there is the belief that nothing short of an absolutely external standpoint—a "God's eye view"—is required for a standard of truth that evaluates our thoughts in light of reality. Third, there is the belief that nothing short of an absolutely external reality—"thing in itself"—is required for such a standard. And finally, there is the belief that nothing short of an absolute proof of the possibility of a strong standard of truth—a proof that is absolutely immune to criticism, both in the present and in the future—is required for surviving the skeptic's attack. It is only by renouncing these illusory requirements that we can make any progress in substantive philosophical theorizing, both in the field of truth and in other philosophical fields.

To prepare the ground for our next principle of truth, the correspondence principle (or, rather, our own version of this principle), let us consider another skeptical argument against truth, the so-called *diallelon* argument, which is directed at this principle.

IV. *The diallelon*

The diallelon is one of the oldest and most influential arguments against the possibility of a standard of truth, usually directed against the correspondence theory of truth. According to the *Merriam-Webster Dictionary*, “diallelon”, or “diallëlon”, is a Greek word, “διαλλήλη”, the neuter form of “diallëlos”. Other related terms are “diallelus” and “dialexis”. “Diallelon” means “definition in a circle”; “diallelus” means “reasoning in a circle”. The notion has also been used for definitions and reasoning involving infinite regress. Since our theory is holistic, it is immune to arguments based on circularity or infinite regress per se. But one of the diallelon arguments, the so-called “comparison” argument, goes beyond circularity (infinite regress) per se, and is particularly important in connection with truth.

The comparison argument essentially says that our thoughts cannot be compared to *reality* but only to *other thoughts*, and therefore, a satisfiable standard of “measuring up to reality” for thoughts is impossible. We can measure one thought against another, but not against reality, or, more generally, against anything other than a thought. The comparison argument has been advanced, or at least taken seriously, by many philosophers, past and present. According to McDermid (1998), these include Spinoza, Berkeley, Hume, Kant, Hegel, Bradley, Brentano, Santayana, Heidegger, James, Dewey, Neurath, Hempel, Blanshard, Goodman, Quine, Rorty, Putnam, Kuhn, Davidson, and Habermas, among others. Some names in this list might be questioned, but clearly the comparison argument has had a broad appeal.

A classical formulation of the comparison argument appears in the Jaesche version of Kant’s logic lectures:

Truth, it is said, consists in the agreement of cognition with its object[:] ... my cognition, to count as true, is supposed to agree with its object. Now I can compare the object with my cognition, however, only *by cognizing it*. Hence my cognition is supposed to confirm itself, which is far short of being sufficient for truth. For since the object is outside me, the cognition in me, all I can ever pass judgment on is whether my cognition of the object agrees with my cognition of the object. The ancients called such a circle in explanation a *diallelon* (Kant 1770s–1800: 557–8).²⁵

This formulation emphasizes both the alleged circularity and the alleged impossibility of comparing thought and reality. More recent versions of the comparison argument say:

- (a) *It is impossible to compare things of different kinds.*

[T]he notion of comparing our system of beliefs with unconceptualized reality to see if they match makes no sense (Putnam 1981: 130).²⁶

²⁵ See also Kant (1770s–1800: 61 (Blomberg Logic) and 280 (Vienna Logic)).

²⁶ This is the antecedent of a conditional, but it is an antecedent that sums up Putnam’s view of the matter.

[W]e can only compare languages... with one another, not with something beyond language called "fact" (Rorty 1989: 20).

- (b) *A comparison between thought and reality requires a "God's eye view"—a viewpoint completely external to, and independent of, the human mind. Since such a standpoint is not available to humans, a comparison between human thought and reality is impossible:*

No... confrontation [between beliefs and the world] makes sense, for of course we can't get outside our skins to find out what is causing the internal happening of which we are aware (Davidson 1981: 144).

[T]here is no independent test of the accuracy of correspondence... unless we can attain what [Putnam] calls a God's-eye standpoint—one which has somehow broken out of our language and our beliefs and tested them against something known without their aid. But we have no idea what it would be like to be at that standpoint (Rorty 1991: 6).

- (c) *A comparison between thought and reality must appeal to a "thing in itself". In requiring a comparison between thought and reality sufficient for truth we require a comparison between thought and an absolute, undescrivable reality, completely independent of the human mind and inaccessible to it—a "thing in itself". This requirement is "incoherent":*

[T]he whole idea of comparing our conceptual system with a world of things-in-themselves... to see if the conceptual system 'copies' the unconceptualized reality is incoherent (Putnam 1983b: 177).

We have already responded to (b) and (c) in our discussion of the fundamental principle of truth, and in particular, the transcendent mode of thought. Comparing a given thought to reality requires no more than moving to a *human* standpoint above (or outside) it, a standpoint like that of a Tarskian metalanguage, which allows us to view both the thought in question and those aspects of the world it relates to. There is nothing godly or superhuman about transcendence of this kind. Similarly: truth is a human standard, a standard designed by and for humans, and as such it is designed to compare human thought with those aspects of the world it relates to—aspects that are in principle open to human view rather than aspects that are in principle closed to it, like a Kantian thing-in-itself.

As for (a), there are a few pertinent criticisms of this argument in the philosophical literature. Van Cleve, for example, makes the sound point that the diallelon confuses the activity of comparison with what is compared, i.e., it confuses cognition with the cognized, representation with the represented:

[The diallelon] argument is not unanswerable. It is no doubt true that I can "compare" a given cognition with its object only by using another cognition, but that is not to say that *what* I am comparing is cognition 1 with cognition 2. The argument confuses the vehicle of knowledge with its content. All one need assume, in order to maintain that it is possible after all to compare a cognition with its object (or as I would prefer to say, to confirm a cognition by checking its object), is that some cognitions do disclose features of the objects they cognize.

Putting the point another way, if the diallelus argument is to reach its conclusion, it must be assumed that cognition *never* discloses or gives knowledge of features of the object cognized, but only of its own self. When the needed assumption is thus made explicit, I doubt that many will want to embrace it. . . . The problem of the diallelus, I would therefore say, arises from a false presupposition (Van Cleve 1999: 216).

The mistake of confusing the external world (which our representations represent) with our representations of it is also pointed out by Kitcher. Speaking about perception, he says:

Perception is a process in which we perceive objects that are typically independent of us by *being in* (or *having*) representational states. We do not perceive by *perceiving* our states (Kitcher 2001: 15).

And speaking more generally, he says:

[W]e should not confuse the possibility of constructing representations with that of constructing the world [i.e., what is represented by our representations] (Kitcher: 51).²⁷

The fundamental principle of truth sharpens our understanding of the diallelon by situating it in the realm of immanent and transcendent thoughts. If, simplifying, we represent the domain of human thoughts and its objects in a cumulative hierarchy of levels:

- Level 0: Things in the world other than thoughts,
 - Level 1: Things of level 0 + immanent thoughts and other mental activities (e.g., perception) directed toward things of level 0,
 - Level 2: Things of level 0 and 1 + immanent thoughts and other mental activities directed at things of level 1: thoughts that attribute truth to thoughts of level 1, mental comparisons of things of levels 0 and 1, etc.,
 - Level 3: Things of level 0–2 + immanent thoughts and other mental activities directed at things of level 2,
- ∴,

then the diallelon thought, like other skeptical thoughts, first arises on level 3. Thus, on level 0 there are objects such as ice (pieces of ice) and properties such as temperature. On level 1 there are *objectual thoughts* about some of these things (e.g., “Ice is cold”). On level 2 there are *truth*, *attribution*, and *comparison* thoughts concerning things of level 1 and their connections to things of level 0 (e.g., “‘Ice is cold’ is true/not true”, or “‘Ice is cold’ attributes to ice the property of being cold”, “Ice has the property that ‘Ice is cold’ attributes to it”). And on level 3 there are various *philosophical thoughts* about truth like the *diallelon*.

²⁷ Pertinent criticisms of the “comparison” argument are also given by Alston (1979), Davidson (1988), McDermid (1998), and others.

What does the diallelon at level 3 say? It says that thoughts of level 2 cannot be true because one cannot compare level 1 thoughts with level 0 things. This is a version of (a) above. Let us look at these claims more closely. The diallelon claim we are concerned with is:

- (D1) *We can only compare a thought with another thought, not with an object which is not a thought.*

Question: Why can we only compare a thought with another thought? One possible answer is:

- (D2) *We can only compare X with things of the same type as X.*

D2, however, is problematic. First, it does not specify what kind of “type” is involved. Clearly, for some types of “type” D2 is utterly implausible. Consider a photograph of a person. If being a photograph is a type, then D2 implies that I can only compare a photograph of my son with another photograph of my son. But even a diallelonist would say that I can compare a photograph of my son to something of another type, for example, my *mental representation* of my son. Second, we *can* significantly compare any one thing with another *so long as we identify a significant point of comparison*.²⁸ In the case of truth, we compare *attributed properties* with *possessed properties* with respect to their sameness as *properties*.²⁹ We ask whether the properties attributed to a given object by a given thought are possessed by that object.

Another possible answer to the question we posed to D1 is:

- (D3) *Human cognition has no access to non-cognitive objects.*

D3, however, is highly problematic. It is implausible that human mental acts are altogether divorced from the world. It is implausible that one cannot be afraid of a lion but only of one’s image of a lion, that one cannot hope that one’s child will recover from a serious illness but only that one’s image of one’s child will undergo a mental process of becoming an image of a recovered child, and so on. More generally, it is utterly unreasonable that creatures which exist in the world, are affected by the world, have to survive in the world, can use the world to their advantage, and have an intellectual interest in the world, are incapable of mentally relating to the world. In the absence of conclusive evidence to the contrary, this view should be rejected. I conclude that at least in the forms considered here, the diallelon poses no serious threat to truth, to the fundamental principle of truth, or, for that matter, to the correspondence theory of truth.

²⁸ Schlick (1935: 66) goes even further: “It is my humble opinion that we can compare anything to anything if we choose”.

²⁹ See footnotes on my use of “property” and similar terms in Section 8.4.

8.3 The “Manifold” Correspondence Principle

Another core, universal principle of truth concerns the connection between mind and reality inherent in it. Traditionally, this principle is called “correspondence”, but the form it takes in the present theory differs in several ways from its traditional form(s). In particular, our version of correspondence allows it to take multiple forms, including highly intricate forms, in contrast to the single and simple form—copy (picture, mirror) or isomorphism—that most traditional conceptions of this relation demand. Correspondence, for us, is a family of interrelated relations, each connecting mind to reality according to the kind of objects a given thought (statement, theory) targets and the resources available to us for cognitively reaching these objects. To signal both the similarities and the differences between traditional correspondence and the one proposed here, I will call the latter “manifold correspondence”.³⁰ As a starting point, we can briefly formulate the principle of manifold correspondence as follows:

(M-COR) *Truth is a matter of a substantial and systematic connection between thought and reality, a connection that has to do both with the way the world is and the way our mind operates. This connection might be quite intricate and take different forms in different fields. The forms it takes depend both on what aspect of reality a given thought targets and on the cognitive resources available to us for reaching it. Abstracting from differences, this connection holds between a given thought and reality when the aspect of reality it targets is, directly or indirectly, yet systematically, as it says it is.*

As a principle of truth, manifold correspondence adheres to many points made by past and present correspondence theorists. Some of these are:

To say of what is that it is not, or of what is not that it is, is false, while to say of what is that it is, and of what is not that it is not, is true (Aristotle 4 century BC: 1011^b25).

Truth is defined by the conformity of intellect and thing; and hence to know this conformity is to know truth (Aquinas 1265–9: q. 16, a.2).

It takes two to make a truth. Hence (obviously) there can be no criterion of truth in the sense of some feature detectable in the statement itself which will reveal whether it is true or false (Austin 1950: 118 fn).

³⁰ It is important to emphasize that *manifold correspondence* is *not* correspondence without a general account of the nature of the relation. Much of what we say about correspondence below, including its initial characterization (M-COR) below, its distinctive characteristics, the reason we cannot get away from it, its grounding both in the mind and in the world, and the commonalities and differences of our specific conception of correspondence with/from other conceptions, are tantamount to a general account of correspondence as it is understood by our theory. (Its being “manifold” is also one of its general characteristics.)

It is natural to hold that what is distinctive about sentences that are true or false is that they *represent* reality as being a certain way. True ones represent it as it actually is, while false ones represent it as it is not (David 1994: 31).

[A]mong the [alternative] paths...only the...path to the correspondence conception [of truth] is viable...[T]here is no alternative but to think of the truth of a proposition as conferred upon it, in the general case, by its relations to non-propositional reality (Wright 1999: 223).

[T]he basic thought behind...correspondence theories [is]: beliefs are true or false as a result of (a) the structure of their parts; (b) the reference of those parts to reality; and (c) the way reality actually is (Lynch 2004a: 84).

Another tenet that manifold correspondence shares with other correspondence theories is the well-known equivalence- or T-schema. Here, however, we begin to diverge from the common approach. Our preferred formulation of this schema is:

(T) True $\langle \sigma \rangle$ iff σ

where:

- (i) " $\langle \sigma \rangle$ " stands for a metalinguistic expression designating a sentence σ of a(n object-) language L, and
- (ii) " σ " stands for a metalinguistic sentence *stating the objectual truth conditions of* σ .

This is a variant of Tarski's T-schema, and where it differs from Tarski's, and most other formulations, is in (ii). Whereas standard formulations of (ii) favor disquotational instances of (T),³¹ as in:

(T₁) "Snow is white" is true iff snow is white,

our formulation regards non-disquotational instances as equally acceptable, as in:

(T₂) "Snow is white" is true iff snow reflects light of all hues completely and diffusely.

The central feature of (T), as we conceive of it, is the equivalence between a *truth predication*, which is *linguistic*, and an *objectual predication* which is *factual* or *worldly*, where the latter states the conditions under which the former holds. And this is captured equally well by (T₁) and (T₂).

Another distinctive feature of our correspondence principle is its *universality*; i.e., we regard *manifold correspondence* as a universal principle of truth, applicable to all areas of knowledge. This view is shared by some, but not all, correspondence theorists; for example, it is not shared by those who view truth in physics as governed by a correspondence standard and truth in mathematics as governed by a coherence standard. Our motivation for viewing correspondence as universal comes from the

³¹ Possibly under translation.

same views that led us to correspondence in the first place: knowledge, qua knowledge, is knowledge of reality, therefore the standard of truth for all knowledge is a correspondence standard; immanent thoughts, qua immanent thoughts, say something about reality (or some aspect of reality), therefore the truth of all immanent thoughts is measured by reference to reality. An intuitive, common-sense observation in support of the universality of correspondence (jointly with the fundamental principle of truth) proceeds as follows.

There is an important sense in which we can never get away from correspondence truth, no matter what we say. Even when we reject correspondence in favor of some other conception of truth, we cannot get away from correspondence altogether. To see why, suppose that you and I stand in meta-arithmetic, or in meta-meta-arithmetic, and you say: "The truth of arithmetical statements is purely conventional"; i.e., you look at the totality of arithmetic statements and deny that their truth value is based on correspondence. Now, this itself is an immanent claim, a direct factual claim about some subject matter, and I am free to ascend to the next level in the metalinguistic hierarchy and raise the critical question: "Is arithmetic truth in fact conventional?" with respect to it. Of course, you are free to refuse to answer my question (as you are free to refuse to answer any question), but the question arises all the same. Suppose you say: "It is more convenient (simple, efficient, fruitful) to think of arithmetic as conventional than as true-based-on-correspondence". This gives rise to the truth question: "Is it in fact simpler, more efficient, more fruitful, to think of arithmetic truth in this way?", and so on. So long as you make an immanent statement, the question of correspondence arises with respect to your statement. The breadth of correspondence is the result of the ubiquity of immanent thoughts together with our freedom to transcend any immanent thought and raise the critical question of whether it measures up to reality, or to that aspect of reality it is directed at. Transcendent claims, too, are for the most part immanent; in particular truth claims are immanent, and as such subject to the question of correspondence. We may say that whatever level of immanent thought we are at at a given time, the question of truth-as-correspondence arises at the next level of thought, a level that we are always free, and able, to ascend to. This renders the correspondence principle universal.

Our inability to get away from correspondence altogether has been pointed out by other philosophers as well. Thus, Kit Fine says:

[I]t seems to me that [even] the antifactualist... should be willing to acknowledge that there is a *factual* standard of correctness. Of course, [for the anti-factualist] the... standard of correctness will be nonfactual... But this nonfactual standard of correctness lives in the shadow, as it were, of a factual standard. For the correctness of our judgments must somehow engage with the real world; there must be something which we aim for in belief and whose realization is a factual matter. So for the expressivist, for example, the factual standard of correctness for a judgment might be that it faithfully reflects one's (possibly implicit) commitments; while for the mathematical formalist, it could be that the judgement is in accordance with the rules of the game. Thus, even in these cases, there is something of a factual sort that counts as getting things right (Fine 2001: 23).

The same holds for the coherentist, epistemic, and pragmatist conceptions of truth. The questions of correspondence here are whether a sentence *S* in fact *coheres* with our body of knowledge, whether an assertion *A* (or theory *T*) will in fact be *justified* at the end of inquiry, and whether a theory *T* (belief *B*) is in fact *pragmatically advantageous*.³²

Both the fundamental principle of truth and the manifold correspondence principle are unifying principles. They say that truth is (always) a transcendent normative standard for immanent thoughts and that satisfaction of this standard (always) requires a certain type of connection between immanent thoughts and reality. But the manifold correspondence principle captures the diversity of truth as well. And in so doing it diverges from traditional correspondence. Three distinctive features of manifold correspondence are:

- A. *Correspondence is, in a deep sense, a matter of both world and mind.*
- B. *Correspondence need not have a simple form, like mirror, copy, or isomorphism; it might have a far more intricate form.*
- C. *Correspondence might take several forms, varying from field to field.*³³

Our reasons for viewing correspondence as “manifold” in this way have their roots in the diversity of truth on the one hand and the basic human cognitive situation which is partly responsible for it on the other. Given the combination of a complex world and limited cognitive resources, humans could either acquiesce in modest cognitive goals or insist on bold goals, with significant ramifications both for the way they go about cognizing the world and for the standards of truth that are suitable for their theories. It is quite clear that human civilization made the second choice. This means (as we have noted in Chapter 5, Section 1) that in order to pursue our epistemic goals we had to stretch our cognitive resources, devise indirect routes for reaching regions of the world which are not directly accessible to us, improvise, experiment, tinker, exercise our imagination, and so on. In other words, we use a wide variety of means to reach the world cognitively, and some of these are complex, indirect, jagged. What this implies for our theory of truth is that it has to allow *multiple* “routes” of correspondence between true cognitions and reality, including routes based largely on intellect (see Chapter 5), as well as highly complex routes, for example, routes that combine intellect and sensory perception, and routes that involve intermediate steps of various kinds (see discussion of mathematical correspondence in Section 8.4).

³² We will return to this point in connection with logical conventionalism in Chapter 9, Section 2.

³³ Terminological note: in what follows I will freely shift from talk about correspondence as a *standard* and talk about correspondence as a *relation*. The two are closely related: to satisfy the correspondence standard a given immanent thought must stand in the correspondence relation to the relevant aspect of reality, and standing in this relation is tantamount to satisfying the correspondence standard.

It also means that a simple, one-dimensional, unchanging standard of correspondence will not do. Such a standard will paralyze our creativity, limit us to simple theories of the world, and leave little room for innovation and development. Rather than imposing a correspondence standard suited to everyday discourse—discourse about observable, mid-size physical objects—on all other discourse, we ought to devise a flexible standard, one that takes into account the kind of routes from mind to reality that is available, or optimal, for a given discipline. At issue is whether the *pattern* of correspondence in one area can differ from that in another area. Suppose that in physics reference (an important constituent of correspondence) is largely based on *direct causal relations* between physical expressions and objects in the world. Does this imply that reference in mathematics must also be based on direct causal relations between mathematical language and the world? Suppose the existence of *physical individuals* is necessary for the truth of physical statements. Does this imply that the existence of *mathematical individuals* is also necessary for the truth of mathematical statements? The basic cognitive situation suggests that the answer to these questions is negative. What type of correspondence is appropriate for a given field is always an open question, something to be found out (upon investigation) rather than determined in advance.

And it is here that mind enters into correspondence as well. Since we have no choice but to cognize the world through the prism of our present cognitive machinery, our standard of truth has to take this machinery into account. This machinery involves a wide spectrum of elements, from sensory perception to intellect, including conceptualization, categorization, abstraction, generalization, reflection, combinatorics, analysis, model building, experiment design, and so on. All these shape our thoughts about the world, and with them, the correspondence conditions that have to hold for our thoughts to be true. Had our cognitive machinery been different—had another sensory modality replaced our modality of sight, for example, or had we been incapable of mathematical thought—the kind of conditions that would have to hold in the world for our thoughts to be true would have been different too.

The correspondence standard of truth is, thus, both world-dependent and mind-dependent. Or, focusing on the mind, it is both mind-dependent and mind-independent. It is mind-dependent in being a standard created by and for humans. It is mind-independent, first in being a standard of correspondence with *reality*, and second, in being a standard of correspondence that, for theories whose subject matter is not affected by the existence of humans (say, gravity), disregards human existence.

Recognition of truth's dual roots in mind and reality, however, can be easily subverted by conceptual confusions. One of these is the "all or nothing" pattern of thinking. In the case of truth, the false dilemma is:

Either truth is based exclusively on the world, in complete independence of the mind, or truth is not based on correspondence at all.

This dilemma is usually implicit rather than explicit in discussions of truth, but this does not diminish its distorting force. There is no reason why truth could not *significantly (and appropriately)* depend on the world without being *exclusively* dependent on it. So long as our standard of truth requires true thoughts to “measure up to reality”, it is a *correspondence* standard, regardless of whether other factors, having to do with the mind, are in some way involved. Of course, we could reserve the word “correspondence” for another standard. But the debate on truth and correspondence is not terminological.

Let us conclude our general discussion of the differences between our conception of correspondence and the traditional conception by spelling out some of our objections to one of its paradigms—the copy (mirror, picture) conception of truth:³⁴

- (a) A true thought does not copy or mirror anything—it is not a copy or a mirror image of anything. A true, or more generally, a truth-apt thought *says* something about some subject matter or *attributes* some property to some object. Copies and mirror images are not in the business of *saying* anything or *attributing* anything to anything. As such they are essentially different from truths.³⁵
- (b) A copy or a mirror image is generally *not selective* while true thoughts are. A mirror image of X mirrors everything within a given spatiotemporal region containing X: X’s shape, color, and size, objects in the vicinity of X, etc. In contrast, a true thought—indeed, any thought—about X is generally very selective, attributing a limited array of properties to X and disregarding all other things about (and surrounding) it.³⁶
- (c) The range of thoughts that can, potentially, “copy” or “mirror” reality³⁷ is quite limited. It is easy to see that, and how, observational thoughts could stand in this relation, but it is far from clear that, or how, abstract, metaphorical, and other types of thoughts could. This unnecessarily limits the scope of truth (theories that do not “copy” reality cannot possibly be true), or else the universality of correspondence, hence the unity of truth (truth in some fields is based on *altogether* different principle from truth in others).³⁸
- (d) The idea of a copy or a mirror image suggests a *simple* and a *single* “route of correspondence”, the same for all fields. This, too, unreasonably limits the theory of truth.

³⁴ For additional grounds for rejecting this conception see Goodman (1976).

³⁵ The same applies to representation. (It is OK to talk about true thoughts as representing reality in some contexts, but it is important to realize that strictly speaking they are not. To say something about X, to attribute a property to X, is not the same thing as to represent X.)

³⁶ In this respect the metaphor of a *map* is superior to that of a mirror image. See e.g., Kitcher (2001).

³⁷ Even in the informal way sanctioned earlier.

³⁸ This point is related to the disanalogy between truth and portraiture noted in Chapter 7, Section 2, fn. 27.

Of course, if we treat “mirror image” or “copy” non-literally, these problems disappear, but then the metaphor loses much of its utility, and the actual practice of traditional correspondence theorists—excluding abstract thoughts from correspondence, associating mathematical correspondence with Platonism—are difficult to explain. To see our non-traditional version of correspondence—manifold correspondence—at work, let us turn to mathematics. There we will make use of the new options offered by our theory to devise new solutions to well-known problems with correspondence in a notoriously recalcitrant field.

8.4 Application to Mathematics (A New Theory of Mathematical Truth)

The attempt to develop an adequate account of mathematical truth has come up against serious difficulties, largely due to the abstract nature of mathematics and its objects³⁹. These difficulties have led some philosophers to question whether mathematical theories are true at all, proposing fictionalist or other error-theoretic accounts of mathematics instead. My goal in this section is to utilize the resources developed in this essay to approach the problem of mathematical truth in a new way. These resources enable us to combine a robust view of the connection between mathematical theories and reality with a flexible approach to the structure of this connection. But this by itself is not sufficient for an account of mathematical truth. What we need, in addition, is an understanding of something basic about mathematics itself, something that will enable us to figure out what form truth (correspondence) takes in this field.

In setting out to investigate a widely studied subject matter, we need to give some thought to our starting point. My starting point will diverge from the prevalent approaches to truth in mathematics in a few ways. First, discussions of mathematical truth often center on relatively small units of cognition, such as sentences (“ $7+5=12$ ”) and terms (“7”). In contrast, I will start by considering the *discipline of mathematics as a whole*. Second, I will start with *extremely basic questions*, questions that, albeit clearly relevant, are not often chosen as the starting point of contemporary investigations: “Is there anything in the world for mathematics to be true or false of?”, “Is there anything in the world such that, if we want to attain knowledge of the world, we

³⁹ In this section I use a few terms (such as “object”, “feature”, “property”, “level”, “law”, “formal”, “modal force”, “order”, “predicate”) whose philosophical usage has sometimes given rise to intractable problems. To prevent confusions and misunderstandings I would like to indicate that my use of these terms here and throughout this volume is limited to their familiar, common-sense usage. (I hope to discuss some of the problems arising from their term-of-art, philosophical usage elsewhere.) When these terms first appear in this section I will sometimes remind the reader of this point and offer a brief, informal explanation of my usage in a footnote. Readers might find some of these notes useful later on as well, especially when reading Chapter 10. The first of these terms is “object”.

Object: I use “object” as a general term for things—individuals, properties, relations, functions, and structures of these.

need to develop a discipline like mathematics to study it?”. Third, when thinking about mathematical truth, and in particular mathematical correspondence, it is common to focus on mathematical *objects*, where the paradigm of a mathematical object is a numerical individual⁴⁰, say the individual 7. In contrast, I will focus on the reality of mathematical *features*⁴¹, where the paradigm of a mathematical feature is a second-level cardinality property, e.g., the property SEVEN.⁴² Thus, I will ask the above questions in the form: “Do objects and properties (in general) have *mathematical features*?” or “Do they have features whose study requires a discipline like mathematics?”, rather than “Are there mathematical objects (paradigmatically individuals, like the number 7) in the world?”.

There are several advantages to this starting point. First, by starting with mathematics as a discipline rather than with individual mathematical sentences we achieve a broad perspective, a perspective from which it is natural to inquire about such things as the role of mathematics (the discipline) in our system of knowledge. Then, if we find good reasons to believe that our system of knowledge requires a mathematical discipline for providing, not just *auxiliary mathematical instruments* but also, and primarily, *genuine mathematical knowledge*, this will give us a strong reason to expect that truth is, in principle, applicable in mathematics. Second, by raising basic questions concerning mathematical truth anew we can avoid past biases on the one hand and make full use of newly gained perspectives and/or advances on the other. Thus, if, for example, we come to see that objects and/or properties of objects in the world have mathematical features, this will strongly suggest that mathematical truth is grounded in reality, i.e., that mathematical truth is based on correspondence (of one kind or another) rather than, say, on mere convention.

Finally, by focusing on the reality of mathematical features, rather than on the reality of mathematical individuals, we provide a wider basis for mathematical correspondence than it commonly has. Specifically, we avoid problems concerning

⁴⁰ *Individual*: I use “individual” for an *atomic* object, one that does not have (or is treated, in a given context, as not having) an internal structure (constituents, arguments, etc.) and is not a property. Individuals are commonly denoted by “individual constants”, and more generally, by “singular terms”.

⁴¹ *Feature*: I use “feature” as a loose term for property, relation, characteristic, trait, attribute, quality, etc.

⁴² (i) *Property*: “Property” is understood in its everyday sense. Often, when I talk about properties I mean “properties and/or relations”, and sometimes “properties, relations, and/or functions”. The reader will be able to figure out from the context which usage is appropriate. Both individuals and properties have properties. I sometimes talk about the “behavior” of objects instead of their “properties”, as in “laws governing the behavior (properties) of objects”. For the purpose of the present discussion there is no need to distinguish between properties as intensions and extensions, and I will often identify a property with the extension of an appropriate predicate or with the set of objects having this property (in a given universe of discourse).

(ii) *Level*: I use a simple, familiar, and informal typology of objects and expressions. Individual constants and individuals are of level 0, predicates of individuals and properties (relations, functions) of individuals are of level 1, predicates and properties of properties of level 1 are of level 2, etc.

(iii) I use numerals and lower-case letters for names of mathematical individuals and small capital letters for names of second-level mathematical, or formal, properties.

the existence of mathematical individuals, problems that have contributed to philosophers' disillusionment with the idea of mathematical correspondence. Furthermore, we are able to communicate with those who, for one reason or another, choose to limit their ontology to nominalistically acceptable individuals. If even such individuals and their properties have mathematical features, then even nominalists have a reason to view mathematical truth as based on certain significant connections between language (thought) and reality.

Some, however, might object to our starting question, "Does reality have mathematical features?", on the ground that this question is too broad or too straightforward or too difficult—that it is naive to think that we can give a conclusive answer to this question at the outset of our investigation. But what we are looking for is a way of beginning the investigation—a *fruitful starting point*, not a conclusive final answer. We are looking for a set of compelling common-sense observations that will introduce fresh considerations into the investigation (or remind us of some fruitful, yet possibly neglected, considerations).

Another objection to our approach might focus on the philosophical nature of the investigation. Some might say that the question whether reality has mathematical features is empirical, something for physics rather than for philosophy to deal with. I will not discuss the *precise* division of labor between physics, philosophy, and mathematics here. But as Feynman (1965: 56) rightly observed, the physicist is not interested in knowing the mathematical features of reality in *complete generality*. Nor is he equipped to discover the *general laws*⁴³ governing mathematical features. Similarly, the mathematician is not interested (at least professionally) in the problem of characterizing the route of correspondence from mathematical theories to reality, which is not a mathematical problem, while the philosopher is. Needless to say, philosophy, on our holistic view, does not isolate itself from either science or mathematics in approaching this problem.

I. Mathematics and the world

In explaining our conception of *reality* in Chapter 5, Section 2, we began with the common-sense observation that objects in the world have features of various kinds, and asked whether these include *abstract* features. Our answer to this question was positive, and as an example of abstract features we considered *formal* features⁴⁴. Now, these features are of a kind that in the present context we consider to be

⁴³ *Law*: I use "law" in the everyday sense of the word: "physical laws", "mathematical laws", etc. In some fields, it is more natural to talk about "regularities", and these, too, fall under my notion of "law". I distinguish laws from accidental generalizations by their greater modal force in the sense of holding over more counterfactual situations (where "more" is understood in terms of inclusion). Usually, I focus on laws governing objects in the world rather than on statements in the language categorized as "laws". When I do talk about the latter, I think of them as representing putative laws in the world.

⁴⁴ *Formal*: I use "formal" to characterize a feature of objects (or reality) or a law governing objects (or reality) of a kind that will be characterized in the main text. (A more detailed characterization will be given in Chapter 10, Section 4.)

*mathematical*⁴⁵; therefore our reasoning in that chapter applies here as well. Briefly, we proceeded as follows:

Take any everyday objects, say humans, or human individuals. These are non-controversial objects, accepted by both empiricists and rationalists, nominalists and Platonists. First, we observe that each individual human—say, a student in a given class—has the formal property of being identical to himself and different from other students. Next, we observe that properties of human individuals also have formal, or mathematical, features. For example, the first-level property of being a student in the class has a second-level cardinality feature or property, say, SEVENTEEN. Proceeding to binary relations of humans, we observe that these, too, have some formal properties. For example, the relation of studying-in-the-same-class-as among students has the formal properties of being reflexive, symmetric, and non-transitive. Similarly, properties of humans can be combined, or operated on, by various formal operations: complement, union, intersection, Cartesian product, etc. For example, by applying the formal operation of intersection to the properties of being a student and being a female we obtain the intersective property of being a female student.

We conclude that if humans are real, and if they have first-level properties and relations like those mentioned above, and if these properties and relations have the second-level properties noted above, then objects in the world have formal, or mathematical, features, and such features are real. In the absence of any reasonable challenge to the existence of humans, to their having first-level properties such as self-identity, and to their first-level properties having second-level properties such as cardinality, the above considerations strongly suggest the reality of mathematical features.

But if mathematical features (properties) are real, then we, humans, want to study them. We want to know about them. We want our system of knowledge to include a branch, or branches, of knowledge that study these features. Furthermore, these (real) features, like other (real) features (e.g., physical features) have the potential of exhibiting regularities and being governed by laws. And this potential is realized by formal properties: self-identity is a law, cardinalities are governed by various laws of order, and so on. Moreover, if formal features are governed by regularities and laws, then it is sensible to presume that these regularities and laws, like many other regularities and all laws, have a certain modal force⁴⁶, a modal force that goes beyond their application to actual objects (objects that actually exist, or are instantiated, in

⁴⁵ *Mathematical*: I use “mathematical” to indicate “belonging to, or studied by, mathematics”. In some contexts, however, “formal” and “mathematical” can be used interchangeably, since, as will be explained in the main text, a formal object is an object of the kind that can be studied by mathematics, and a mathematical object (of the kind considered here) is usually a formal object (or a representative of such an object), studied by mathematics.

⁴⁶ *Modal force*: I use “modal force” as a general term of our language, not as a philosophical term-of-art associated with a particular semantics, say, Kripke’s. “Modal force” refers to such things as impossibility, possibility, contingency, and necessity, where each of these can be of several types (e.g., physical vs. formal necessity), and there is a sense of order attached to these notions (e.g., weaker and stronger necessity).

the world). So, even if we start as nominalists with respect to individuals, it is difficult for us to deny that the individuals we sanction and their properties have formal features (e.g., self-identity, cardinality), or that these formal features exhibit regularities and are governed by laws (laws of identity, laws of cardinality, etc.).⁴⁷

II. *The subject matter of mathematics*

The question arises: which discipline studies the formal? That is, which discipline studies the laws governing formal or mathematical features of objects in the world? Since laws have a modal force that goes beyond their actual instantiations, a discipline that studies them must have rich resources. And if, as we will explain below (and in greater detail in Chapter 10), formal laws have an especially strong modal force, then such a discipline must have especially rich resources. Is there a suitable candidate for such a discipline? The natural answer is: “Yes. *Mathematics*”. But does mathematics fit this job description?

Some would give a negative answer to this question. They might say that mathematics is purely conventional, that its job is to set forth formal conventions based on pragmatic considerations, and that if it studies anything, what it studies are just such conventions. Or they might say that mathematics is too general and too abstract to engage with the real world, that it is interested in the formal for its own sake, whether it is real (part of reality) or not. Mathematical conventionalism, however, is highly problematic (for reasons we will discuss in Chapter 9, Section 2)⁴⁸, and in the absence of a compelling defense of this doctrine, it is sensible to pursue other options. As for the claim that mathematics is interested in the formal independently of reality, first I would like to note that mathematics is a broad and diverse discipline, with multiple goals and interests. The question is whether *one* of the central things it does is study the laws governing the formal features of reality, not whether this is the only thing it does.⁴⁹ Second, the view that mathematics is not interested in reality is unreasonable. Consider cardinality. We have seen that properties of individuals in the world have cardinalities. We know that mathematics has theories of cardinality (e.g., arithmetic and set theory). Since these theories can be unified, we can talk about *the* mathematical theory of cardinality. The question is: does, or at least ought, the mathematical

As I progress in this chapter, and especially in Chapter 10, I will make “modal force” more specific and explain what it does and does not mean.

⁴⁷ At this point some readers may wonder how our account deals with the super-large ontology of mathematics or want to see in greater detail how it explains the strong modal force of mathematics. The former will be explained in Subsection IV; the latter will follow from the detailed explanation of the strong modal force of logic given in Chapter 10.

⁴⁸ Although this discussion will deal with logic, its conversion to mathematics is straightforward.

⁴⁹ Indeed, the claim is not even that the only way the formal is integrated in the world is through properties which are purely formal, like cardinality properties. We do not rule out properties which have both physical and formal elements, but these are not our target here. (Hence, “formal”, in this essay, indicates what elsewhere might be called “purely formal”. A more precise notion of “formal” will be given in Chapter 10, Section 4. See mention of “invariance” later in this chapter.)

theory of cardinality encompass “real” cardinalities (i.e., cardinalities of actual or potential objects in the world)? Should its laws (what it says are laws of cardinalities) be real laws of real cardinalities?

Now, a negative answer to this question is quite implausible. It is implausible that properties in the world have cardinalities, these cardinalities are governed by formal laws, mathematicians know about these cardinalities and believe they are governed by laws, yet they study exclusively other laws, laws governing *unreal*, made-up cardinalities—laws that have nothing to do with cardinalities in the world. It would be strange if mathematics studied finite cardinalities, but the laws it attributed to such cardinalities were (knowingly) different from the laws governing real cardinalities. It would be strange if in saying “ $1+1=2$ ” mathematicians had in mind some unreal cardinalities, cardinalities that had no systematic connection with the cardinalities of properties of actual and potential objects in the world. And it is irrational to say that this is what mathematicians ought to do.⁵⁰

It will not do to say that only *applied* mathematics is concerned with *real* cardinalities. First, applied mathematics requires non-applied mathematics in order to apply it to the world; i.e., it requires a general mathematical theory that is not limited to, say, the number of objects that happen to exist in the world at a given time. Second, the task of formulating the laws governing *real* cardinalities in a *completely general* and *systematic* way requires something on the order of “pure” mathematics. This is because of the special nature of *formal laws*. Let me explain. We have noted that laws in general have a wider scope than the actual things over which they range: laws are *counterfactual* in nature. This holds of physical laws as much as of formal laws. But formal laws differ from physical laws in dealing with properties that have a *considerably greater degree of “invariance”*. The exact meaning of this will be explained in Chapter 10, Sections 4–5, but roughly, the idea is that formal properties are indifferent to many features of the individuals in the underlying domain, more features than physical properties are indifferent to. For example, if the underlying ontology includes both physical and non-physical individuals, formal properties will not distinguish between them, but physical properties will. For that reason the laws of formal properties are more general and abstract than those of physical properties, and this means that we need more general and abstract theories to account for them in full generality. In fact, we need theories on the same level of abstractness and generality as “pure” mathematical theories.

This is also the reason mathematics is not limited to cardinalities that are actually instantiated in the world. Suppose we start with finite cardinalities and we want to state in *complete generality* the laws governing these cardinalities. To achieve this goal we cannot make do with a theory that allows a *largest* finite cardinality: we need a theory that does something on the order of recognizing (all) infinitely many finite

⁵⁰ This line of reasoning has affinities with both Frege and Russell.

cardinalities; i.e., the task of accounting for the finite in complete generality brings us to infinite collections of objects (or something like it). But once we recognize infinite collections of objects, we want to know the formal laws governing them (beyond those of finite cardinalities), and this leads us to recognize still larger infinite collections and, with them, infinite cardinalities, indeed an unending series of higher and higher infinite cardinalities. The desire to know the laws governing infinite collections of objects is largely motivated by our intellectual interests; but intellectual interests play a prominent, if not an exclusive, role in our epistemic life.⁵¹ For instance, to state, in sufficient generality (given our interests), the law governing the relation between the cardinality of a given property (considered extensionally) and the cardinality of its “power-property”, we need something as general as the full-fledged “Cantor’s theorem”.⁵² Similarly, to know the ordinal properties of objects, or structures of objects, in the world (in the desired degree of generality), we have to go beyond finite ordinals. Not only do we have to countenance higher and higher infinite ordinals (or something like that), but once we go beyond finite ordinals we need to distinguish ordinal properties and their laws from cardinal properties and their laws. That is, in order to know the formal laws governing reality in a way that is *commensurate* with their *degree of invariance*, we need something on the order of “pure” mathematics.⁵³

Now, if mathematical theories (or some mathematical theories, or significant parts thereof) are theories of the laws governing the formal features, or formal behavior, of objects in the world, then these theories (or significant parts of these theories) are true or false in the *correspondence* sense, broadly understood. If, and to the extent that, the laws of, say, our current arithmetic theory are—directly or indirectly, but in a systematic manner—the laws governing the relations between finite cardinalities in the world, then current arithmetic is true to reality. In other words, if, and to the extent that, current arithmetic or set theory is committed to a correct description of the laws governing real cardinalities, then its standard of truth is a *correspondence* standard. This, however, does not mean that its correspondence to reality is *direct*. It is an open question *how* (in what manner, exemplifying what pattern) true mathematical theories do, or might, correspond to reality, *what* the route of

⁵¹ Notwithstanding the initial unwelcoming attitude toward, say, Cantor’s studies of the infinite, eventually the interest in his theory and its later developments has largely prevailed.

⁵² Cantor’s theorem says that for any set/property X of any cardinality, finite or infinite, the cardinality of $\text{Power}(X)$ is larger than the cardinality of X (where $\text{Power}(X)$ is the set/property of all subsets/subproperties of X).

⁵³ What about the higher reaches of set theory? Our example explains how we arrive at indenumerable cardinals accessible through the power-set operation (given an infinite set of posited individuals) but not the possibility of larger cardinals. We leave the question of larger cardinals open, at least for now. It might turn out that we need such cardinals in order to express some formal laws governing reality in complete generality; we may have good pragmatic reasons for introducing such cardinals; we may discover good reasons, of one kind or another, for rejecting such cardinals; we may continue to be uncertain with regard to such cardinals; etc. All these possibilities are compatible with our approach.

mathematical correspondence is or might be like. Our next task is to figure out what this route is.

III. A “composite” route of mathematical correspondence

In trying to figure out what the route of mathematical correspondence is, an apparent incongruity presents itself. Our analysis suggests that the level at which cardinalities, say, arise in reality is the level of properties of properties, but in modern arithmetic and set theory cardinalities are individuals. This creates a puzzle: why, if the formal (largely) resides on the level of properties of properties, does mathematics study it on the level of individuals and their properties? Why does mathematics study cardinalities, which are second-level properties, by means of first-order theories⁵⁴, such as Peano Arithmetic and ZFC, which construe them as individuals?

The puzzle has two sides. From one perspective the question is: “If cardinalities are in fact second-level properties, how can first-order arithmetic and set theory get things right? How can an arithmetic theory which treats cardinalities as individuals be said to *correspond to reality*?” From the opposite perspective the puzzle is: if, in the world, there are no cardinal individuals but only second-level cardinal properties, why do mathematicians construct their theories of cardinalities (such as Peano Arithmetic and ZFC) as theories of individuals? Why do they treat cardinalities as individuals if in fact they are properties of properties?⁵⁵

Turning to the first perspective, if we put the question in the form, “How can mathematical statements about individuals be true if there are no mathematical individuals?”, the question seems difficult to answer. But if we put it in the form, “How can laws of individual numbers correspond to laws of second-level cardinality properties?”, we immediately see how to answer it. In the same way that a skyscraper built of steel, concrete, and glass can be accurately represented by a small plastic model, so higher-level objects and their laws can be accurately represented by lower-level objects and their laws. First-order theories are capable, in principle, of correctly describing second-level objects and their laws, if we allow indirect representation; first-order statements about individual numbers can correspond to facts and laws involving second-level properties, if we allow *indirect* or *composite correspondence*.

⁵⁴ *Order*: “Order” here is a correlate of “level” (see fn. 42(ii) in this chapter). For example: A first-order theory is a theory whose non-logical constants and variables are of level/order ≤ 1 .

⁵⁵ (i) For the present purpose it does not matter whether mathematicians always prefer first-order theories to higher-order theories. It is sufficient that they sometimes do, or even that they could do so.

(ii) Even if there are mathematical individuals in the world, say, numbers, the problem does not disappear. There are still cardinality properties of properties in the world, so mathematical theories of cardinality have to account for them as well. One might say that second-order arithmetic accounts for these cardinalities. But then one has to explain why there is such a close (if not complete) correlation between cardinality properties and individual cardinals.

But why would humans use composite correspondence in studying the formal rather than simple correspondence? Why not opt for higher-order mathematics?

One way of answering this question is to note that we, humans, seek to know things about the world which might not be directly accessible to us. How do we go about it? That depends on how our cognitive apparatus, and in particular our intellect, works. Suppose we are so constituted that our intellect works more effectively when we deal with systems of individuals than with systems of higher-level objects. Suppose the most natural or effective way for us to make discoveries and develop theories of any subject matter (or of formal subject matters, or of certain formal subject matters) is to do so on the first-level; we may be better at discovering formal regularities and constructing systematic theories of such regularities when we think of them as involving individuals rather than higher-level properties. Our cognitive tools—especially our intellectual tools—may work better in a first-level setting than in a higher-level setting.⁵⁶ Now, if reality itself does not supply us with a first-level formal setting, we can exercise *epistemic freedom* and create a first-level setting for studying the formal by ourselves. We can use our imagination, creativity, ability to discover (or figure out) relations between phenomena, etc., to construct a first-level model (in the everyday sense of “model”) of reality, or of those parts/aspects of reality we wish to study. This enables us to develop a first-order arithmetic theory that gives a correct, albeit indirect, account of cardinalities in the world. Arithmetic, in that case, describes the laws governing cardinalities by describing laws governing their first-level simulations. The key idea is that while the subject matter of mathematics (or of significant parts of mathematics) is external—mathematics seeks to discover formal laws governing *reality*—mathematics is a *discipline created by and for humans*. As such it might reach the world in ways that are advantageous for humans but circuitous from the point of view of correspondence. Such a correspondence can be as *accurate* and *systematic* as direct correspondence in spite of its roundabout nature. Laws of arithmetic and set theory may not be connected to reality through the same route as laws of other sciences, yet they might be connected to it just as deeply.⁵⁷

How would humans go about constructing first-order theories of higher-level cardinalities? One way they might go about it is to introduce a *postulated* level of *individual cardinals* which are systematically connected to *second-level cardinality properties* in the world. Mathematical reference and correspondence would then be *composite*. Using ordinal subscripts to indicate the order of “legs” in such reference and correspondence, we may say that individual numerical terms would refer₁ to posited individual numbers, and refer₂ to second-level cardinality properties which are systematically represented by their referents₁. Similarly, first-order statements expressing laws of cardinal individuals would correspond₁ to posited first-level

⁵⁶ Or, with respect to some tasks we may work better in a low-level setting, with respect to others, in higher-level settings.

⁵⁷ This account bears some similarities to Hofweber (2005: §4).

cardinality laws, and correspond₂ to higher-level cardinality laws which their correspondents₁ systematically represent. The reference and correspondence at work in mathematics will then be reference₁₋₂ and correspondence₁₋₂—two-stage reference and correspondence connecting 0- and first-order expressions with second- and third-level objects.⁵⁸

Figuratively, we can represent composite reference and correspondence as follows:

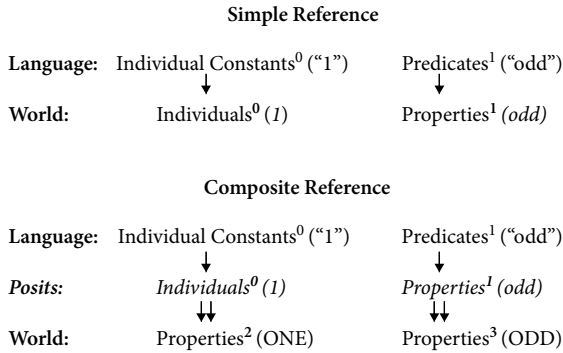


Figure 8.1

Traditional correspondence theory disregards the possibility of using our cognitive resources to reach reality in this way. As a result, it limits us to a single, monolithic pattern of correspondence, the same for all first-order truths; and this, in turn, pressures us to renounce correspondence altogether whenever the standard pattern fails to work. But by taking into account the resourcefulness, flexibility, conceptual capacities, and imagination of humans we see at once that they—we—can change the traditional pattern of correspondence and adapt it to our needs. And one type of adaptation is postulation of a new class of objects or an auxiliary stratum of reality. Two examples of *composite correspondence* based on this approach are:

Composite Correspondence

Example 1

"2 + 7 = 9" is true

iff

+(2,7)=9

iff

DISJOINT-UNION (TWO, SEVEN) IS NINE⁵⁹

⁵⁸ This would be an example of the intellectual activity of "figuring out", exercised either consciously or unconsciously.

⁵⁹ Where disjoint union is the union of two properties whose extensions are pairwise disjoint.

Example 2

“($\forall n$)($\forall m$)($n+m = m+n$)” is true
iff
the operation of addition is symmetric
iff
DISJOINT-UNION IS SYMMETRIC.

A systematic connection between the first- and second-level truth conditions of mathematical statements, however, does not mean reduction or translatability to an *equivalent* theory.⁶⁰ Although we can translate first-order number statements to second-order number statements, first-order arithmetic is *not equivalent* to second-order arithmetic: first-order arithmetic has a logically complete proof system (the proof system of standard first-order logic) while second-order arithmetic does not; second-order arithmetic is categorical while first-order arithmetic is not. The two are not equivalent, yet they are systematically connected.⁶¹

Speaking objectually, the posited stratum of mathematical individuals need not be *fully* reducible to a non-posed layer of reality. In a sense, the posited stratum has a *life of its own*. Once the basic structure of a lower-level model of the formal features of reality is in place, the road is open for the lay as well as the theoretical mathematician to utilize it in studying these features and, in particular, their laws. That is, once the representational adequacy (in principle) of the posited stratum is accepted, we proceed *as if* mathematical truth were based on simple (direct) correspondence with this posited stratum. (It was thus by taking seriously the picture of sets as individuals that mathematicians made a quantum leap in their understanding of the hierarchy of sets—a hierarchy involving sets of higher and higher levels.)⁶²

Studying the formal layer of reality indirectly in this way generates certain discrepancies, and it is our duty to recognize them, find ways to reconcile them, and see to it that no paradoxes are generated. But with the availability of moves like *transcendence* (ascent to a metatheory or a sideways step to a background theory), back-and-forth movement between center and periphery, shifts in position on Neurath’s boat, and multiple holistic foundational tools, this can in principle be done. We can identify with considerable precision the relation between first- and second-order arithmetic, we can structure the network of connections between posited cardinality individuals and cardinality properties of different levels in a sufficiently intricate way to prevent paradox, and so on. All these make a composite connection between mathematics and the world not just a viable option but also a

⁶⁰ Indeed, the viability of a theory indirectly connected to reality does not require the existence of a counterpart which is directly related to reality.

⁶¹ It is misleading to say that second-order mathematics is a worldly mathematics while first-order mathematics is a human mathematics. Both have human as well as worldly facets, although their “routes” to reality are different.

⁶² Our account of numbers as representing higher-order cardinality properties bears some similarities to Hodes (1984, 1990) and Hofweber (2005, referred to earlier), as well as to Yi (1998, 1999) and Moltmann (2013).

feasible one. Responding to the skeptic we will say that *if you know how to connect them to reality, the laws of mathematics do not lie*.⁶³

IV. Merits of the present account

Although our general claim that mathematical truth is based on correspondence is in principle compatible with more than one account of mathematical correspondence, the present account has several advantages. First, unlike deflationist accounts of mathematical correspondence,⁶⁴ it is a substantive, explanatory account. Second, unlike most substantivist correspondence accounts, it is neither an empiricist nor a Platonist account, and as such it is free from the familiar problems with mathematical Platonism and mathematical empiricism. Third, it has considerable problem-solving capabilities, especially when combined with our holistic approach to truth and knowledge. Fourth, it has the ability to make a significant contribution to the foundations of logic. The first two advantages should be clear by now; the fourth will become clear in Chapter 10. Here I will focus on the third.

Composite correspondence (together with holism) enables us to solve, or at least make significant progress toward solving, a wide range of problems in the philosophy of mathematics. These include the identity problem, the applicability problem, the modal force problem, the large-ontology problem, the cognitive-access problem, and the mathematics-as-algebra problem, a few of which are pertinent to mathematical realism and correspondence. A full discussion of these problems would be exceedingly long, but briefly we can indicate the problems and their solution, or progress toward solution, as follows:

(a) *The identity problem* (Benacerraf 1965). Benacerraf's problem is the problem of the identity of mathematical individuals. There are many distinct, yet isomorphic (standard) models of first-order arithmetic, including models in which the same numerals are assigned different referents. For example, in Zermelo's model the numeral 2 is assigned one object—the set $\{\{\emptyset\}\}$ —while in von Neumann's model it is assigned a different object—the set $\{\emptyset, \{\emptyset\}\}$.⁶⁵ Which object is the real 2?

On the present account the problem does not arise. Both Zermelo's and von Neumann's numbers 2 are posits representing the second-level property TWO, and since TWO is a formal property, all isomorphic systems of posits for it are equally good. That is, Zermelo's 2 represents the second-level TWO just as well as von Neumann's 2 (in their respective systems). Posited individuals need to have some identity, but what identity we give them is immaterial, so long as the result is a systematic representation of the target properties.

This, indeed, explains why some philosophers are inclined to view mathematics as conventional. The choice between one adequate system of posits and another is

⁶³ This way of formulating the response is inspired by the title of Cartwright (1983).

⁶⁴ "Deflationist"—in the sense discussed in Chapter 7, Section 1.

⁶⁵ Here " \emptyset " names the empty set, and " $\{x, y\}$ " names the set of x and y .

largely conventional. But on a more fundamental level mathematics is not conventional: the system of numerical posits represents formal features in the world. At its core, arithmetic is a theory of real cardinalities (the finite ones), and the laws governing these cardinalities are, for the most part, not conventional. Since the laws of arithmetic represent these laws, they, too, are not conventional.

Should we say, then, that the “real” 2 is the higher-level property TWO?⁶⁶ From the point of view of our account of first-order mathematical truth above, the answer is “Yes”. But from a transcendent perspective which views our account as a tentative proposal, the question is open.⁶⁷

(b) *The applicability problem.* The applicability of mathematics to the world, including those aspects of the world that are studied by empirical disciplines, is a crucial aspect of mathematics’ role in our system of knowledge. Yet due to mathematics’ highly abstract nature, its applicability to the empirical world is traditionally considered difficult to explain.⁶⁸

[T]he enormous usefulness of mathematics in the natural sciences is something bordering on the mysterious[;] ... there is no rational explanation for it (Wigner 1960: 2).⁶⁹

⁶⁶ I use “higher-level” rather than “second-level” here because cardinality properties can be of a higher level than 2. For example, cardinality properties of second-level entities are of a higher level. A more comprehensive account of numbers as cardinality properties will generalize to cardinality properties of any level ≥ 2 .

⁶⁷ Our solution to the identity problem suggests a lesson for a more general, largely neglected, problem concerning truth and correspondence: *the truth-of-a-sentence vs. truth-of-a-theory problem*. In the philosophical literature, truth is commonly viewed as a property of sentences. But epistemically, the truth of theories is more significant. Not much has been written about the relation between truth of sentences and truth of theories, and the default view appears to be that a theory is true iff all its sentences are true. Our discussion of the identity problem above shows, however, that this approach is far too simplistic. Suppose you have two theories, T1 and T2, such that for some sentence S , $S \in T1$ and $\sim S \in T2$. Is it possible that both T1 and T2 are true in the correspondence sense, i.e., both T1 and T2 correspond to reality? Intuitively, the answer is “No”, but in fact, this depends on the circumstances. Let T1 be Zermelo arithmetic, let T2 be von Neumann arithmetic, and let S be the statement “Successor $\{O\} = \{\{O\}\}$ ”. Then S is true in T1 and false in T2; hence, according to the intuitive answer, T1 and T2 cannot both be true. Yet T1 and T2 are both accurate arithmetic theories, so if one of them is true, so must the other be. Our account captures the sense in which there is no real conflict between the two theories: T1 and T2 represent the system of laws governing the behavior of cardinalities in two different, yet equally adequate ways. A “Tarskian” analysis would explain the situation by appealing to translation: T1 is a translation of T2 to a Zermelo language (T2 is a translation of T1 to a von Neumann language). But our account suggests that direct translation is not needed. The relation between the truth of a sentence and the truth of a theory can be more complex than the Tarskian analysis suggests, yet still comprehensible. An especially important case-study is that of truth across scientific revolutions (in Kuhn’s sense), where direct translation is presumably not available. I will leave the discussion of this case for another occasion.

⁶⁸ For **boldface** within citations please see fn. 3, Chapter 1.

⁶⁹ Wigner’s view is open to a variety of interpretations. Steiner (1998: 9), for example, regards Wigner as speaking “not of discovery, but of description”, and he interprets Wigner as talking about “individual success[es]” of mathematical applications rather than about a “successful... grand strategy”. The intended interpretation of Wigner’s statement in the present context is different: his statement, as it is read here, has to do both with discovery and with description, and is concerned with a “grand strategy” rather than with “isolated act[s]” of application. In thinking about the applicability of mathematics to science in terms of a general strategy, our approach is similar to Steiner’s.

The problem is to explain how abstract mathematical laws apply to non-abstract objects, phenomena, situations. This problem is especially difficult for traditional, Platonist adherents of mathematical correspondence. Since Platonists believe in the existence of two disconnected realities, one abstract, the other physical, they must explain how mathematical laws which govern a Platonic reality apply to a radically different, mundane physical reality. But this is an exceedingly difficult task.

This problem does not arise for our proposal. First, our proposal assumes only *one reality*. Second, it connects mathematics with formal *features* of objects, not with the existence of mathematical individuals. Third, it holds that (our single) reality has both physical and formal features. And fourth, it says that formal features are in principle features of objects of various kinds, including physical objects. As a result, the applicability of mathematics to physics is quite straightforward. The question is not how to connect mathematical individuals to physical individuals; the question is how to apply formal laws to physical objects (individuals and properties). But this is not problematic. Physical objects have properties of many kinds, including formal properties, and therefore laws governing such properties apply to them in an unproblematic way. Indeed, the abstractness of the formal makes it easier, rather than more difficult, to apply it to non-formal objects. The relevant trait of formal features is that they *abstract from*—i.e., *disregard*—most features of their bearers, and this means that they do not “notice” the difference between their physical and abstract bearers. In the language of Chapter 10, they have an exceptionally high degree of invariance. For example, cardinality features “treat” collections of physical objects just like collections of abstract objects, and for that reason they apply equally, and in the same way, to both.⁷⁰

(c) *The modal force problem.* If mathematical truth is truth about the world, then it is presumably contingent; but mathematical truths are necessary rather than contingent (see e.g., Thalos 2013). On our account, this is a false problem. Mathematical truth is truth about the laws governing formal features of objects in the world, and it depends on the nature of formality whether these laws are necessary or contingent. A precise explanation, in terms of invariance, will be given in Chapter 10.⁷¹ Briefly, the higher the degree of invariance of given features of objects (any objects, including physical objects), the larger the number of counterfactual situations in which objects have these features, hence the larger the number of counterfactual situations in which their laws (non-vacuously) hold. Since the degree of invariance of mathematical features is especially high, they disregard more features of objects (situations) than more weakly invariant properties (e.g., physical properties), and therefore the counterfactual range of their laws surpasses that of physical laws. But the range of

⁷⁰ More generally: When laws are abstract in the sense of abstracting from, or disregarding, particular cases, their abstractness is an advantage rather than obstacle for broad applicability.

⁷¹ Readers unfamiliar with my notion of invariance may delay reading the rest of this paragraph until reading Chapter 10, Section 4.

counterfactual situations in which a law (or a true statement) holds determines its modal force; hence mathematical laws have an especially strong modal force. In short, the formality of mathematics explains its considerable modal force.

(d) *The large-ontology problem.* It is sometimes thought that mathematical realism and traditional correspondence have difficulty explaining the large ontology of mathematics. There appears to be a gap between the immense ontology of mathematical individuals and the presumably far smaller ontology of “real” individuals. The problem is to justify mathematics’ claim to such a super-large ontology.

On our account, the problem does not arise. So long as the sizeable collection of mathematical individuals is a collection of *posited* individuals, its size, by itself, poses no genuine ontological problem. If, and to the extent that, a large layer of posited entities is required (or advantageous) for a precise, informative theory of the laws governing the formal features of objects in the world, positing such a layer is in principle warranted.

Indeed, our proposal does more than dissolve the large ontology problem. It explains why mathematics needs such a large collection of posited individuals. Mathematics, as a theory of the formal, is a theory of laws governing features of objects. But laws in general, as we have noted above, have a certain degree of invariance, hence necessity, and as such require a counterfactual ontology. Formal laws have an especially strong degree of invariance; hence they require an especially large counterfactual ontology (or something like it). Earlier, we gave the example of Cantor’s theorem to clarify this point: to express the laws of the natural numbers in full generality we need something on the order of a denumerable ontology of individuals. And once we use a denumerable ontology of individuals, we need something on the order of an indenumerable ontology of individuals to express laws governing denumerable collections of individuals.⁷²

(e) *The cognitive-access problem (Benacerraf 1973).* The problem of cognitive access is the problem of how to access those aspects of reality that are responsible for the truth or falsehood of mathematical statements and theories. As formulated by Benacerraf, the problem is:

Given that our only access to reality is *causal*, how can we have cognitive access to mathematical *objects* (e.g., the numerical individuals 0,1,2,...), which do not stand in any causal relation to us?⁷³

Benacerraf’s problem involves two assumptions: (i) humans’ *only* cognitive access to reality is causal, and (ii) mathematical correspondence requires the existence of

⁷² It should be noted that our account does not determine the size of the universe of mathematical individuals uniquely. Many questions are left for mathematics itself, metamathematics, and the philosophy of mathematics to consider, and various answers to these questions are compatible with our account.

⁷³ Benacerraf is also concerned about a semantic-syntactic parity in delineating the route of mathematical correspondence. This requirement, which we have already challenged, is left aside here.

mathematical objects (as distinct from mathematical properties), and in particular mathematical individuals. In addition, (iii) Benacerraf does not consider the possibility of non-traditional correspondence (for example, correspondence involving posits), assuming the only possible route of correspondence is the simple traditional route. Finally, he seems to assume that (iv) there are only two options for pursuing mathematical realism: the traditional options of Platonism and empiricism. Given all these assumptions, Benacerraf's problem is indeed resistant to solution.

None of these assumptions, however, holds for our account. The cognitive-access problem, as viewed from our perspective, is the problem of how we, humans, have cognitive access to *mathematical features* of reality and the laws governing them, given their abstractness, generality, and strong modal force. A solution to this problem requires extensive investigation, on a number of levels and by a number of disciplines (mathematics, psychology, philosophy, and possibly others). But as far as the philosophical component of this investigation goes, our account of knowledge and truth points to a rich trove of cognitive resources, including resources suitable for cognizing abstract features and general laws. The combination of a holistic yet world-oriented conception of knowledge, a dynamic view of our system of knowledge (disciplines moving from center to periphery and vice versa), basic realism, multiple forms of cooperation between intellect and sensory perception, multiple routes of correspondence, and so on, opens the door to a new conception of mathematical knowledge. This conception allows mathematical knowledge to rely on both causal and non-causal connections between reality and mind,⁷⁴ circumvents problems concerning mathematical individuals,⁷⁵ exploits epistemic freedom in managing our cognitive access to abstract features of reality, and grounds the modal force of mathematical laws in the strong invariance of formal features. The account as a whole locates formal (mathematical) features and their laws right here, in the actual world, so the issue of access to a separate, Platonic, reality does not arise.

(f) *The "mathematics as algebra" problem.* The inner diversity of mathematics has led some philosophers to distinguish between *algebraic* and *non-algebraic* mathematical theories. Shapiro characterizes the two by reference to their intended subject matter. Non-algebraic theories are "about a single structure that is unique up to isomorphism" (Shapiro 1997: 40), algebraic theories are not. Arithmetic is a paradigm of a non-algebraic theory; group theory of an algebraic theory. In model-theoretic terms:

A field is nonalgebraic if it has a single "intended" interpretation among its possible models or, more precisely, if all its "intended" models are isomorphic (or at least equivalent). . . . A field is algebraic if it has a broad class of (nonequivalent) models (Shapiro: 50).

⁷⁴ E.g., connections involved in "figuring out" (see Chapter 5, Section 3).

⁷⁵ E.g., through focusing on mathematical features rather than mathematical individuals and construing mathematical individuals as posits.

Our account of mathematical truth so far has focused on theories that Shapiro would characterize as non-algebraic. What about algebraic theories? Algebraic theories, on our account, relate to the world in two ways: First, they are *conditionally* true or false, second, they potentially *model* certain aspects of reality. Viewed in the first way, algebraic theories are mathematical theories in the “if-thenist” sense (Putnam 1967b). They are true if, and to the extent that, their theorems follow logically from their axioms (assuming their set of axioms is logically consistent). Viewed in the second way, algebraic theories can, in principle, model various aspects of (or structures in) the world, and when they do, they are said to be *true of* those aspects (structures).⁷⁶

V. *Relation to other philosophical views*

Our conception of mathematical correspondence, and manifold correspondence more generally, has interesting relationships with other conceptions. I will not be able to trace all these relations here (or, indeed, any one of them with great detail), but I will briefly discuss points of similarity and difference with a few: Aristotle, Frege, Quine, some contemporary fictionalists, and Horgan.⁷⁷

(a) *Aristotle*. Our conception of mathematics bears some similarities to Aristotle’s, especially as construed by Lear (1982). Reality, according to Aristotle, has multiple aspects, and this is reflected in the variety of features possessed by physical objects. In particular, physical objects have mathematical features (like being spherical), and it is these features that are the subject matter of mathematics. Mathematics, therefore, studies *real* features of *real* objects, and its theories are genuinely true (or false). How does mathematics study these features? By separating them (in thought) first from the physical objects that possess them and second, from other features these objects possess, so it can study them on their own. Thus, Aristotle says:

Obviously physical bodies contain surfaces, volumes, lines, and points, and these are the subject matter of mathematics....[T]he mathematician...separates them, for in thought

⁷⁶ It is worth noting that we see no sharp difference between algebraic and non-algebraic theories. It is common to view arithmetic as a *theory* of a certain aspect of the world and group theory as a *model* of certain phenomena in the world, but we can view arithmetic as providing a *model* of, say, counting, and group theory as a *theory* of, say, coordinate transformations between inertial frames. Models, like theories, can account for a given set of phenomena in the world either accurately or inaccurately, either systematically or not systematically, and theories, as we have seen above, can involve the postulation of entities just like models. Furthermore, our dynamic holism sanctions fluctuations in status between “being a theory” and “being a model”. For example, it enables us to view Euclidean geometry as having turned from a “theory” (of geometrical space in general, physical and abstract) into an “if-thenist” theory, or a possible “model” (of some regions of physical space), and to view Riemannian geometry as having turned from an “if-thenist” theory or a “model” into a “theory” (of the geometrical structure of physical space).

⁷⁷ Significant relations also exist with Russell (1903/38, 1919), Whitehead and Russell (1910–3/1925–7), Hellman (1989), Resnik (1997, Chapter 9), Fine (2005), Sider (2011), the authors noted in fn. 62 in this chapter, and others. See brief comments on some of these in fn. 88 in this chapter.

they are separable from motion . . . and it makes no difference nor does any falsity result if they are separated (Aristotle, *Physics* 193b23–35. Quoted in Lear 1982: 162).

This effective way of studying mathematical features might involve the positing of mathematical entities:

The best way of studying [such features] would be this: to separate and posit what is not separate [i.e., what is not separate in reality but is separable in thought], as the arithmetician does and the geometer (Aristotle, *Metaphysics* 1078a21–3. Quoted in Lear 1982: 165. Material in square brackets inserted by me).

Under this conception, mathematical objects exist, but they exist in a special way: namely, as abstractions from physical objects. This relation between mathematics and physics explains how mathematics applies to physical objects. Lear concludes:

For Aristotle, mathematics is true, not in virtue of the existence of separated mathematical objects to which its terms refer, but because it accurately describes the structural properties and relations which actual physical objects do have. . . . [There is no need] to explain mathematical truth . . . via the existence of mathematical objects. One can understand how mathematics can be true . . . by understanding how it is applicable (Lear 1982: 191).⁷⁸

These citations point to several significant similarities: Both accounts start from the observation that objects in the world have formal/mathematical features and that one central task of mathematics is to study these features; both sanction the postulation of mathematical objects that represent such features; both regard mathematical truth as genuine truth; both regard mathematical truth as based on some kind of circuitous correspondence (Aristotle implicitly; we explicitly); both allow cognitive access to formal/mathematical features of reality; and both have resources for explaining the application of mathematics to physics in a relatively straightforward manner.

But beyond these citations there are significant differences between the two accounts. First, our own account is not committed either to Aristotle's special ontological theory of real, non-positated, objects, or to his views on the variety of formal features of objects in the world. Second, our account does not require that a property be instantiated in order to be real. Third, our account explicitly talks about laws governing the behavior of formal features of reality, and it sanctions laws that are farther reaching than those Aristotle had in mind (for example, various set-theoretical laws that were not contemplated by Aristotle). Fourth, our account has resources (e.g., Fregean resources), that were not available to Aristotle, for dealing with arithmetic truth, as well as resources (such as invariance) for explaining the necessity of mathematical laws. Furthermore, there are methodological differences

⁷⁸ This account (and explanation) applies smoothly to geometry, but Aristotle uses a somewhat different (and arguably weaker) account for arithmetic. For us, however, it is natural to extend Aristotle's account of geometry to arithmetic, by viewing numbers as representing cardinality properties of physical objects.

between our approach and Aristotle's. For example, to the extent that Aristotle's methodology is foundationalist, our own methodology is holistic. As such, it provides us with a wide variety of tools for grounding mathematical theories in reality and explaining mathematical truth/knowledge that were not available to Aristotle.

(b) *Frege*. Manifestly, there are significant similarities between our theory and Frege's. Our conception of cardinalities as second-level properties, our focus on systematic connections between cardinality properties and numbers as individuals, our view of mathematics as objective, our emphasis on a close relation between logic and mathematics (see Chapter 10, Section 8)—all these have some roots in Frege. But there are also important differences.

One difference concerns the relation between language and ontology. In spite of saying that "it is one of the tasks of philosophy to break the domination of the word over the human spirit" (Frege 1879: 7), Frege gives more weight to language as a key to ontology than we do. For him, humans' use of singular terms to indicate numbers is a sign that numbers are individuals. (See e.g., Frege 1884.) For us, the direct relation between language and ontology is far less binding.⁷⁹ Natural language is a multipurpose tool, shaped by evolution, psychology, sociology, history, and chance. It was developed in a variety of contexts, and often in a messy, haphazard way. As a result, there is no reason to assume a precise correlation between its terms and things in the world. Theoretical language is less messy and haphazard, but it, too, cannot be taken as a guide for ontology. It is greatly influenced and partly shaped by natural language, earlier theories, and a variety of accidental circumstances. As such it is sometimes less than ideal as a measure of ontology, especially in such fields as mathematics, where access to reality involves a different balance between intellect and sensory perception than in other areas. Furthermore, theoretical language has multiple goals, including simplicity, efficiency, ease of communication, etc., and these too might stand in some tension with the goal of ontological transparency. We must also mention human creativity, including the capacity to create "fictional" posits (systematically related to non-fictional objects), build (artificial) models of certain aspects of reality, and so on. Indeed, not only do we have these capacities, we have good reasons to exercise them. This is something we might do both instinctively and deliberately, both haphazardly and in a planned manner. As such, it would be manifested both in natural and in theoretical languages. Under these circumstances, drawing rigid conclusions about the nature of reality from the surface structure of our language is problematic.⁸⁰

⁷⁹ This was noted in Chapter 5.

⁸⁰ It should be noted, however, that although Frege is usually thought to have proceeded from the linguistic level of number terms to the reality of numbers as individuals, there are interpretations of Frege, or at least accounts inspired by Frege, that construe his individual numbers as posited entities. Hodes (1984), which we have already mentioned, falls under this category. Hodes regards the positing of numerical individuals as a notational "coding device" and the number theorist who posits them as a "coding-fictionalist". Among the advantages of such a *coding* (compared with second-order discourse) are

Another difference from Frege concerns the status of mathematical truths. While Frege characterizes mathematical truths as analytic, we characterize them as factual. And while Frege views them as purely apriori, we, as holists, view them as “quasi-apriori”, i.e., as arrived at largely through the use of intellect, but not in complete isolation from experience (see Chapter 5, Section 3). Still another difference concerns logic. Although we share Frege’s view that mathematics and logic are closely related, we differ on the precise nature and structure of this relation (see Chapter 10, Section 8). Finally, we reject Frege’s claim (in e.g., Frege 1918) that truth is primitive and unexplainable. Our account of truth as composite correspondence is based on a substantivist, non-primitivist approach to truth.⁸¹

(c) *Quine*. The idea that positing objects is central to human knowledge is a Quinean idea. Not just mathematical objects but also abstract physical objects, everyday common-sense objects, and even subconceptual sense data are posits, according to Quine—theoretical, conceptual, and evidential, respectively. What is the point of positing such diverse types of object? For Quine, the point is pragmatic. In the case of molecular particles, “the particles are posited for the sake of a simple physics” (Quine 1955: 250). Mathematical objects are posited to expedite scientific knowledge in general, and they are justified by their *indispensability* to empirical science. What can we conclude from the fact that objects of all types are posits?—“[T]hat posits are not *ipso facto* unreal” (Quine: 251). In fact, posits are essential for our very notion of reality: “it is by reference to [everyday bodies, which are posits] that the very notions of reality and evidence are acquired” (Quine: 252). The boundary between the real and the posited disappears.

Our own approach is similar to Quine’s in its emphasis on the basic and positive role of posits in knowledge, but is different on other counts. First, we are not committed to the view that all objects are posits. Second, we think of posits as playing a broader (and more fundamental) role in knowledge than a purely pragmatic role. This is because posits, on our view, enable us not just to simplify our theories but also to do more fundamental things like *discover* certain aspects of reality that we might not be able to discover without them (think of discovery as involving “figuring out”, as discussed in Chapter 5, Section 3). And posits may also illuminate the structure of certain theories. Third, while Quine regards mathematics as subsidiary to empirical science, we see it as standing on its own. Mathematics’ primary role is to provide knowledge of the laws governing formal features of reality, and its role in physics, though very important, is secondary to this role. Fourth, while Quine recommends a criterion of ontological commitment that relies on the *syntactic* features of a given

a less “messy” notation, smaller logical complexity, greater familiarity (Hodes: 144–5), and citing Quine (1969c: 112–13), “an extraordinary combination of depth and simplicity, beauty and utility”.

⁸¹ Here, too, however, we recognize the existence of divergent interpretations of Frege. Bar-Elli (2010), for example, offers an interpretation of Frege’s conceptions of analyticity and apriority under which our approach is closer to Frege than under more common interpretations.

theory, we believe his criterion is one of merely “surface” ontological commitment. Syntax is one of the least flexible features of language and one of the last to catch up with new discoveries and developments concerning the “makeup” of the world. As such, its role in registering the true ontological commitments of our theories is rather limited.⁸²

(d) *Contemporary fictionalists*. Although our account offers an *alternative* to the fictionalist account of mathematical truth, it also bears some similarities to it. Contemporary fictionalists, such as Field (1989), recognize that theories need not refer to “real” objects in order to be “good” theories; they appreciate the importance of human creativity in knowledge; and they recognize the advantage of *fictional* objects over “real” objects in avoiding various philosophical difficulties. In all these, our approach resembles theirs.

Furthermore, although Field believes that mathematics plays a purely *pragmatic* or *instrumental* role with respect to science and is *conservative*⁸³ with respect to science, he emphasizes that this does not mean that we can eliminate mathematics through direct translation to a “pure” (nominalistic) scientific language. Nor does it mean that there is a problem with accepting those parts of mathematics that have no application to science, like the higher reaches of set theory. These, Field says, are natural ways of extending the “story” told by the useful parts of mathematics (Field 1989: 10), and as such are legitimate. Field is also flexible with respect to justification. “[J]ustification”, he emphasizes, “is not an all or nothing affair” (Field: 17). And among the methodological principles he appeals to are such broad principles as “inference to the best explanation” (Field: 14–20). Our view is similar to Field’s in these as well as other respects.

But we also differ from Field in some fundamental ways. Field’s basic approach to knowledge is physicalistic, and although he does not ban all abstract objects (e.g., he endorses space-time points), he is a strict nominalist when it comes to mathematics. We, in contrast, do not object *in principle* to the existence of mathematical objects, and we do view reality as having genuine formal features. Most importantly, we do not measure the epistemic value of mathematics only by its instrumental value for physical knowledge, but view it as an independent source of knowledge. As a result we, unlike Field, view mathematical theories as capable of being genuinely true, true in the sense of correspondence, albeit composite correspondence, with reality. Nor do we limit the role of mathematics in empirical science to “conservative extension”. On the contrary, we regard the applications of mathematics in science as applications of genuine formal truths. Finally, our view differs from fictionalism with respect to logic. Mathematical fictionalism puts a heavy burden on logic: the dispensability of

⁸² See also discussions of ontological commitment and the relation between language and ontology in Chapter 5.

⁸³ A mathematical theory *M* is conservative with respect to science iff for any nominalistic scientific theory *N* and a nominalistic sentence *S*, *N*+*M* logically implies *S* only if *N* (by itself) logically implies *S*.

mathematics is due to the fact that all nominalistic results arrived at using mathematics can be derived from other nominalistic results using only logic. Logic itself they regard as “real”, in contrast to mathematics. But logic, for them, is not real in the sense of being anchored in reality: logic is disengaged from reality. In this respect their view of logic is more traditional than ours. While we share their view that logic is real, we regard it as real in just the sense of being anchored in reality, as we explained in Part II and will explain in greater detail in Part IV.

Due to these differences, most current objections to fictionalism do not apply to us. One alleged problem for fictionalism that might be thought to apply to us was raised by Yablo (2001).⁸⁴ We can formulate it as follows: on the one hand the fictionalist is ready to assert that there is an even prime number, hence, that there are numbers; on the other hand, the fictionalist does not believe in the reality of numbers, hence is committed to asserting that there are no numbers. So the fictionalist is both committed to the view that there are numbers and committed to the view that there are no numbers. (Another way to put it is to say that the fictionalist is committed to the self-defeating sentence “The number of numbers is 0” (Yablo 2001: 80).) Yablo dissolves the problem by showing that it results from overlooking a simple distinction: the distinction between “*representational aids*” and “*things-represented*” (Yablo: 81), between “figurative” and “non-figurative” speech (Yablo: 85). The crucial point is that number words “can travel back and forth between the two categories”; indeed “they can do it . . . within a single sentential move” (Yablo: 81).

Yablo’s solution, however, is just in the spirit of our conception of knowledge, language, and composite correspondence. When we say that *there are individual numbers* we refer directly to the posits and indirectly to reality (cardinality properties); when we say that *there are no individual numbers* we refer directly to reality, and not at all to the posits. When we say that the number of numerical individuals is 0 we use “numerical individual” as a term ranging over real individuals and “0” as a term denoting the posit 0, which represents the second-level property ZERO. So there is no contradiction. Our point is simply that in reality cardinals are of level 2 rather than 0, that 0-level cardinals are merely posits, and that these posits might be nonetheless central to human knowledge. Indeed, Yablo’s talk of “travel[ing] back and forth between the two categories” finds its natural home in our dynamic model of knowledge, holistic methodology (Neurath’s boat), and emphasis on the centrality of “transcendence” to knowledge and truth.

(e) *Horgan*. Horgan (2001)⁸⁵ develops a theory of “indirect” correspondence geared toward empiricists. The theory, as Horgan characterizes it, is a combination of metaphysical realism and a liberal correspondence theory of truth, where metaphysical realism is essentially the view that the world consists of a definite totality of

⁸⁴ An earlier version, directed as *modal* fictionalism, is due to Rosen (1990).

⁸⁵ See also Barnard and Horgan (2006, 2013).

discourse-independent objects and properties. Correspondence is the view that truth depends, at least in part, on the way the world is, and liberal correspondence is the view that this dependence need not involve a direct, 1-1 correlation between singular terms and objects in the world, predicates and properties in the world,⁸⁶ quantifiers and sets of objects in the world (those over which they range), etc. The basic idea is that truth depends both on the world, independently of the mind, and on the totality of our semantic standards (falling under mind, in our sense). Those standards vary from one context to another, and as a result there is a whole spectrum of ways in which, and degrees to which, a statement's truth value can depend on the world. At one end of the spectrum are contexts for which the semantic standards require a direct connection between statements and reality; at the other end contexts for which those standards require no more than an *empty* connection with reality, i.e., *no* connection with reality at all. The intermediate zone consists of contexts and statements for which the semantic standards require indirect connections with reality of various types, levels, and degrees.

Horgan's conception of correspondence is empiricist or naturalistic in nature. Its empiricism manifests itself on three levels. (a) The real ontology of the world is empirical: there are no abstract objects or properties. (b) The relevant relations between language and the world, as far as truth is concerned, are exclusively empirical—specifically, causal. (c) The null end of correspondence is large and philosophically consequential. In particular, all pure mathematical truths as well as moral truths are based on *null* correspondence.

As an example of a truth based on indirect correspondence Horgan proposes:

- (1) Beethoven's fifth symphony has four movements (Horgan 2001: 73).

The truth of (1), Horgan (2001: 73) argues, "does not require that there be some ENTITY answering to the term 'Beethoven's fifth symphony' and also answering to the predicate 'has four movements'." Rather, it is sufficient that there are "other, more indirect, connections between [(1)] and THE WORLD", connections with, in particular, Beethoven's and others' "musical" *behavior*. This includes Beethoven's "composing his fifth symphony", his earlier compositional activities ("in virtue of which his later behavior counts as composing his *fifth* symphony"), "a broad range of human practices (including the use of handwritten or printed scores to guide orchestral performances) in virtue of which such behavior by Beethoven counts as 'composing a symphony' in the first place", and so on.

By accepting indirect correspondence Horgan is able to considerably expand the range of correspondence truth open to empiricists. Even those who deny the existence of abstract objects can acknowledge that a sentence like (1), which on the surface talks of abstract objects and properties (symphonies, movements of symphonies, the

⁸⁶ *Predicate*: A predicate is an expression that normally indicates, denotes, or stands for a property (see fn. 42(i) in this chapter). An n -place predicate, $n \geq 1$, stands for an n -place property (relation).

property of being fifth in a series, and the property of having four parts), has a definite truth value, and its truth value is to a significant degree a matter of fact—a matter of how things are in the physical world.

One consequence is that two sentences may share the same grammatical structure yet differ in the patterns of correspondence their truth involves. Consider, for example:

- (2) Beethoven's second city of residence has two rivers.⁸⁷

This sentence has a similar syntactic structure to (1), yet it corresponds to reality in a direct (or a more direct) way, according to Horgan's theory. Another consequence is that a single sentence may be connected to reality in a great number of ways, and different connections are relevant to its truth in different contexts.

As for (1) itself, it is linked to reality both directly, through a causal link to a bona fide physical object, Beethoven, as well as indirectly, through causal chains connecting the abstract terms "symphony", "movement", "fifth symphony", and "four movements" to physical events. But (1) is not linked to reality through any *independent* connections of the ordinal term "fifth" and the cardinal term "four" to any features of reality. Ordinal and cardinal term, like any other mathematical terms, are not connected to reality at all. Instead they are connected (exclusively) to pure conventions, i.e., in our terminology, to the mind.

It is here that the radical nature of Horgan's empiricism displays itself. On Horgan's pluralism of forms of correspondence, true statements are linked to reality in two ways: through direct causal links to explicitly mentioned, empirically acceptable objects (individuals, properties, events, etc.), and through indirect causal links to mentioned and unmentioned empirically acceptable objects (individuals, properties, events, and so on). Some of these may be indirectly represented in language by *abstract* singular terms and predicates. But no terms are connected, directly or indirectly, to abstract features of reality, features whose reality Horgan seems to deny.

Our own theory goes far beyond Horgan's. In particular, it extends the idea of indirect correspondence to *all* fields of knowledge, including highly abstract fields such as logic and mathematics. It further shows that, and how, this can be done without accepting highly problematic traditional doctrines like Platonism, which empiricists rightly shun. And it integrates this view into a dynamic model of knowledge and a foundational-holistic methodology which might appeal to moderate empiricists as well as to moderate rationalists.⁸⁸

⁸⁷ This example is mine.

⁸⁸ In fns. 62 and 77 (of the present chapter) we indicated a few other related views. Briefly, our similarities with Russell (1903/38, 1919) are related to realistic trends in his work; we differ from him in ways that enable us to avoid some problems with his account, e.g., concerning infinity. Hodes (1984, 1990) offers an interpretation of Frege that anticipates certain aspects of our composite correspondence. Hellman's (1989) *modalist mathematical realism* does not require a commitment to mathematical

Next, let us use our example of mathematical correspondence to discuss the more general problem of mixed sentences and inferences.

VI. *The problem of mixed sentences and inferences*

Our theory of manifold correspondence might be viewed as a pluralistic correspondence theory. The pattern of correspondence in one field might vary from that in another. This is one of the tools we use to balance the unity and disunity of truth: truth in all fields (to which truth applies) is based on correspondence, but correspondence itself can take multiple forms, depending on features of the world that a given discourse is directed at as well as features of the discourse itself. The question arises how we deal with truths and inferences of mixed kinds. As we recall from Chapter 7, Section 2, this question was raised as a challenge to another pluralistic approach to truth, namely that which regards truth as based on correspondence in some fields and coherence or ideal justification in others. This latter kind of pluralism is, in certain respects, more extreme than ours, and Tappolet (1997, 2000) challenged it to explain how it deals with logically mixed sentences and mixed logical inferences. If the truth of some sentences is of the correspondence type and that of others of the coherence type, what is the type of truth of logical constructs of mixed sentences? Or consider logical inferences involving mixed truths. How can a truth of one type yield a truth of another type? This is especially problematic when some of the (essential) premises have what Tappolet (1997: 209) calls “lightweight” truth (e.g., a coherence standard of truth) while the conclusion has a “heavyweight” truth (e.g., a correspondence standard of truth). How can the coherence truth of one sentence guarantee (or play a crucial role in guaranteeing) the correspondence truth of another?

This problem does not arise for our theory, since it construes truth in all fields as correspondence. But a related question might be asked: what *form* does correspondence take in logically complex sentences whose atomic components take different forms of correspondence, and how can correspondence of one form guarantee correspondence of another form in a mixed logical inference? Although a deeper understanding of the matter must await our discussion of logical inference in Chapter 10, we can already say the following:

Logic, on our view, is subject to a correspondence standard of truth, just like any other field of knowledge, so the question is how the route of logical correspondence/reference (whatever it is) interacts with other routes of correspondence/reference. Consider the logically valid first-order inference,

individuals. Resnik (1997, Chapter 9) suggests that mathematical individuals are posits. And Fine (2005) offers a detailed, clear, elegant, and precise account of the kind of positing we have in mind—“postulating”, in his terminology.

- (3) The number of oxygen atoms in a water molecule is 1
 The number of hydrogen atoms in a water molecule is 2
 $2 > 1$

 The number of hydrogen atoms in a water molecule > the number of oxygen molecules in water molecule.⁸⁹

Assume the route of mathematical correspondence (reference, satisfaction) is *composite*, while that of physical correspondence (reference, satisfaction) is *simple*. Now, since (3) is a first-order inference and logical inference takes into account only the precise reference/satisfaction conditions of the logical vocabulary, (3) takes into account only a schematic first-order version of the reference/satisfaction conditions of the non-logical vocabulary. This means that (3) takes into account *only* the (schematic form) of the *first "leg"*⁹⁰ of the composite route of reference (satisfaction) of its mathematical vocabulary. Since this leg is of the same kind as the route of reference (satisfaction) of the physical vocabulary, there is no problem of incommensurability. Once we arrive at the conclusion, we can go back to the more complex account of reference for the mathematical vocabulary, and this will result in a complex account of its correspondence conditions. More generally, the correspondence conditions of mixed truths will be mixed as well.⁹¹

Our conception of truth as based on manifold correspondence is extendable beyond mathematics to other fields: logic, abstract physics, ethics, etc. We will explain its extension to logic in Chapter 10, while physics and ethics will be left for a future study. Still, let me say a few words about ethics. We have already discussed a few options concerning moral truth earlier in this chapter in connection with the fundamental principle of truth. Our discussion of correspondence suggests a few additional points. In investigating truth in ethics we should not try to fit moral truth into the Tarskian template of a recursive definition of truth, which is largely geared toward logic. The direction of investigation should go from ethics to truth rather than from truth to ethics. To understand truth in ethics what we need to do is what moral philosophers always do: try to understand morality in its own terms. Only then will we be able to find out whether truth operates in this field, and if it does, how. Among other things, we need to ask whether there is something objective in moral judgments, what this objective element is and how it grounds moral judgments, how we, humans, cognitively reach this grounding, whether the connection between moral

⁸⁹ One way to formalize this inference so the issues we are talking about are displayed, is:

$$\frac{(\forall x)(Wx \supset (\iota y)Oyx = 1) \quad (\forall x)(Wx \supset (\iota y)Hyx = 2)}{2 > 1} \quad \frac{}{(\forall x)(Wx \supset (\iota y)Hyx > (\iota y)Oyx),}$$

where " ι " is the definite description operator and the symbolization key for the non-logical vocabulary is: Wx - x is a water molecule, Oxy - x is the number of oxygen atoms in y , and Hxy - x is the number of hydrogen atoms in y .

⁹⁰ See Figure 8.1 in this chapter.

⁹¹ We will expand on this point in Chapter 10.

judgments and their objective ground is one of truth, i.e., correspondence, what pattern of correspondence is involved, how it relates to patterns of correspondence in other fields, and so on.⁹²

8.5 The Logicality Principle (Tarski's Theory of Truth in Perspective)

We have seen two universal principles of truth: *the fundamental principle* and *manifold correspondence*. Both are *core* principles of truth, i.e., principles that capture something very *basic* and *general* at the core of truth, something *central* to *all* truths. The two principles are interconnected: the correspondence principle elaborates on one aspect of the fundamental principle: the normativity of truth. It says that truth as a *standard* for immanent thoughts requires a systematic connection between true thoughts and reality, yet this connection may take different forms in different fields, and to understand the standard of truth in a given field is to know what form it takes in that field. In this way, the principle of manifold correspondence plays a major role in balancing the unity and diversity of truth.

The logicality principle is a third *universal* principle of truth. It says that *logical structure* is a central factor in rendering sentences true or false, a factor that works in the same way in *all* fields of truth. As such, the logicality principle is a *unifying* principle of truth. But it is not a core principle of truth in our sense. It is concerned with a *partial* and *highly specific* determinant of truth, one among many other determinants. Aside from the specific case of *logical truths and falsehoods*, logical structure does not determine, all by itself, the truth value of sentences, but combines with other factors, factors pertaining to physical correspondence, mathematical correspondence, etc., to determine their truth value. And in the case of logically atomic sentences, i.e., sentences containing no logical constants, logical structure plays no role at all.

But although the logicality principle is concerned with a partial and highly specific determinant of truth, this determinant is uniquely important, lying, as it does, at the root of a powerful and systematic method of inference essential for knowledge. Indeed, we can hardly make any significant step in thought and reasoning without introducing logical structure into our discourse, and the difference this structure makes to the truth conditions of our thoughts is substantial. It is important to note that it is the specific features of the logical determinant, those distinguishing it from all other determinants of truth, that render the logicality principle universal, that is, at work in *all* fields of knowledge. Logical structure, by its nature, abstracts from, or disregards, all characteristics of sentences that distinguish those belonging to one field of knowledge from those belonging to another. And this enables logical

⁹² To make this sentence manageable, I left out qualifications such as “if there is” and “if it does”.

structure to affect the truth value of all sentences in the same way in *all* fields.⁹³ The logical factor is, thus, an example of a factor of truth that, while substantive and universal, is partial and highly specialized—an example, in fact, of a factor whose substantiveness and universality are closely related to its partial nature and specialization.⁹⁴

It is important to recognize, however, that the logicity principle, like the fundamental principle and manifold correspondence, is not a one-line principle: its content cannot be fully expressed by a single and simple definition or statement. Rather, its content requires a cluster of statements, definitions, and even theories, concerning the role of logical structure in truth and the role of truth in logic. In studying the logicity principle, therefore, we have several important tasks:

- A. Describe the contribution of logical structure to truth in a way that explains how it gives rise to a powerful method of inference.
- B. Give a well-reasoned, normative account of the logical constituents of human thought (on the linguistic level: logical constants and logical form), and critically examine the current account of these constituents.
- C. Examine whether the logicity principle is grounded exclusively in the mind (language, concepts), as it is traditionally thought to be, or also significantly in the world, as our general considerations in Parts I–II and earlier in the present Part suggest. If the latter, investigate what aspects of the world it is grounded in.
- D. Explain how the logicity principle of truth is related to the other universal principles of truth

And so on.

A critical contribution to the first task was made by Tarski in two monumental papers: “The Concept of Truth in Formalized Languages” (1933) and “On the Concept of Logical Consequence” (1936a). Let us now turn to the first of these.

I. Tarski’s theory as a theory of the contribution of logical structure to truth

In “The Concept of Truth in Formalized Languages” Tarski offers a systematic account of the contribution of logical structure to the truth value of sentences. He does this by constructing a systematic method for defining truth for sentences of given languages, a method centered on the logical complexity of those sentences. Tarski’s method (or definition, as it is usually referred to) has been discussed in many

⁹³ For a detailed discussion, see Chapter 10.

⁹⁴ (i) The fact that logic deals with a partial and highly specific determinant of truth explains why the view, common in the early twentieth century, that logic is the “be all and end all” of philosophy is unwarranted. The fact that logic deals with an extremely important and universal determinant of truth explains why the view that logic is a marginal field of philosophy, quite common today, is equally unwarranted.

(ii) In viewing the principle of logicity as a non-core principle of truth, I diverge from Frege (1918); in viewing it as a universal and highly important principle of truth, I follow him.

works.⁹⁵ Here I will offer an explanatory account of his method which is new in some respects and coincides with other accounts in others. This account will enable us to respond to major criticisms of Tarski's paper, but primarily it is designed to contribute to the theory developed in the present essay.

TARSKI'S METHOD OF DEFINING TRUTH

Tarski's method of defining truth is guided by two conditions: a condition of *material* adequacy and a condition of *formal* adequacy. The condition of material adequacy, as Tarski conceives of it, is intended to ensure that the definition of truth captures the intended notion rather than some other notion.⁹⁶ What is the notion of truth that Tarski seeks to account for? Tarski declares early in his paper that:

throughout this work I shall be concerned exclusively with grasping the intentions which are contained in the so-called *classical* conception of truth ('true—corresponding with reality') (Tarski 1933: 153).

He regards this notion as captured by a "*semantical definition*" of truth (Tarski: 155), where by a "semantical" definition or a semantical notion (concept) he means a definition or a notion that deals with the relation between language and the world. Thus, speaking generally about semantic concepts he says:

A characteristic feature of the **semantical concepts** is that they give expression to certain **relations between the expressions** of language and the **objects** about which these relations speak (Tarski: 252).

And explaining his use of "semantics", he says:

We shall understand by **semantics** the totality of considerations concerning those concepts which, roughly speaking, express certain **connexions between the expressions** of a language and the **objects and states of affairs** referred to by these expressions (Tarski 1936b: 401).

That is, semantic notions are, essentially, correspondence notions, in the broad sense in which "correspondence" is used in our theory. Turning to semantic definitions, Tarski says that a semantic definition of truth is "a definition which we can express in the following words":

- (4) *[A] true sentence is one which says that the state of affairs is so and so, and the state of affairs indeed is so and so* (Tarski 1933: 155).

⁹⁵ See Field (1972), Soames (1999), and Patterson (2012), among many others.

⁹⁶ Tarski's use of "material adequacy" is different from some current uses. Aside from the explanations given in his 1933 paper, he says elsewhere that a materially adequate definition "*grasp[s] the current meaning of the notion as it is known intuitively*" (Tarski 1931: 129), and again that a "materially adequate... definition of truth... aims to catch hold of the actual meaning of an old notion": such a definition will "enable anyone to determine whether the definition actually fulfills its task" (Tarski 1944b: 341).

Tarski's self-designated task in his paper is to make the "intuitive meaning and general intention" of (4) "more definite, and to give it a correct form" (Tarski: 155).

In the case of particular sentences, his example of a semantic definition of truth is:

(5) '[I]t is snowing' is a true sentence if and only if it is snowing (Tarski: 156).

The idea is that (leaving indexicals aside) "It is snowing" is true iff what it says about the world, namely, that it is snowing, is the case, i.e., iff it is snowing.⁹⁷

The condition of formal adequacy has a different purpose. Its task is to ensure that the definition of truth is "in harmony with the laws of logic" (Tarski: 165). Most importantly, its task is to ensure that this definition does not entail a contradiction or lead to a paradox. The danger of paradoxes in the field of truth has been known since antiquity. A sentence that says of itself that it is false—for example:

(6) (6) is false—

is true if false and false if true (assuming bivalence).⁹⁸

Tarski's precise formulation of the formal adequacy condition is based on his solution to the Liar paradox. His solution does two things: (a) It limits the definition of truth to "open" languages, i.e., languages which do not contain either their own semantic predicates (including "true") or names of their own expressions. (b) It relegates the definition of truth for a given language L to an appropriate metalanguage, ML, where ML is stronger than L in several ways. For example: in addition to having resources for saying all the things that can be said by sentences of L, ML has resources for naming all the expressions of L and attributing semantic properties and relations to them, including relations that connect them to objects and properties in the world. Tarski further limits his attention to languages formulated within the framework of modern Fregean-Russellian logic, languages which he calls "formalized languages of the deductive sciences" (Tarski: 153).

⁹⁷ Since sentences such as (5) are instances of the equivalence schema, those philosophers who identify this schema with the disquotational conception of truth might regard Tarski's semantic conception of truth as a disquotational conception, where this is understood to be opposed to the correspondence conception. Tarski himself sees no conflict between the two:

A characteristic feature of the semantical concepts is that they give expression to certain relations between the expressions of language and the objects about which these relations speak, or that by means of such relations they characterize certain classes of expressions or other objects. **We could also say** (making use of the *suppositio materialis*) that these concepts serve to set up the correlation between the names of expressions and the expressions themselves (Tarski 1933: 252).

We can explain Tarski's view as follows: There are two modes of speech, an *objectual mode* and a *linguistic mode* ("material", in medieval terminology). The correspondence idea can be expressed in both modes, and this is reflected in the equivalence schema and its instances, which can be read in either mode. Tarski's overall conception of truth and semantics in the 1933 paper suggests that he gives *philosophical priority* to the objectual, correspondence reading of (5) over its linguistic, disquotational reading. In any case, it is under this reading that Tarski's theory fits into ours.

⁹⁸ Throughout this essay I assume bivalence, but it is important to reemphasize that bivalence is not essential for my theory (see Chapter 5, Section 4, and fn. 38 in Chapter 10, Section 4). Nor is it essential for Tarski, in my view.

Assuming the formal adequacy condition is satisfied, the material adequacy condition is concisely formulated by Tarski as follows:

Given a language L and an appropriate metalanguage ML of L , the definition of the truth predicate of L in ML has to satisfy the following condition:

CONVENTION T. *Given any ML-name, X , of an L-sentence, S , the definition implies the ML-sentence:*

X is true iff S^ ,*

where S^ is an ML-sentence which says the same thing as the L-sentence S .⁹⁹*

(It is readily seen that CONVENTION T captures the idea underlying (5).)

Now, in setting out to formulate his general method for defining truth in accordance with his formal and material conditions of adequacy, Tarski chose to proceed by means of an *example*. This is a point that has generated some misunderstandings, so it requires clarification. Tarski explains his choice as follows:

For an extensive group of formalized languages it is possible to give a **method** by which a correct definition of truth can be constructed **for each of them**. The **general abstract description** of this method and of the languages to which it is applicable would be troublesome and **not at all perspicuous**. I **prefer therefore** to introduce the reader to this method **in another way**. I shall construct a definition of this kind in connection with a particular concrete language and show some of its most important consequences. The indications which I shall then give in §4 of this article will, I hope, be sufficient to show how **the method illustrated by this example can be applied to other languages of similar logical construction** (Tarski: 167–8),

It is readily seen that Tarski intended to develop a *general method* for defining truth, i.e., a method for defining truth for all (or for a broad array of) formalized languages, and his motivation for presenting his method by means of a specific example was merely pragmatic, having no essential philosophical significance. This will also be seen from the consequences he drew from his definition and the main use he made of it:

Tarski's example. Using contemporary terminology, Tarski's example can be concisely described as follows:

- (i) *Object Language: L_C .* The target language is the language of the calculus of classes (a theory with one non-logical constant, the inclusion relation between classes). I will refer to it as " L_C ". L_C is an interpreted first-order language whose primitive vocabulary consists of the logical constants " \sim " (negation), " \vee " (disjunction) and " \forall " (the universal quantifier), the non-logical constant " \subseteq " (a 2-place predicate interpreted as class inclusion), and variables, " x_1 ", " x_2 ", " x_3 ", ..., ranging over all objects in the domain, D_C , of L_C , where D_C is a set of classes.

⁹⁹ Based on Tarski (1933: 187–8) (but limited to the material adequacy part).

- (ii) *Metalanguage: ML_C* . Truth for L_C is defined in a metalanguage, ML_C . ML_C relates to L_C in the way described above. In particular: (i) the syntax of L_C is describable in ML_C , (ii) each expression of L_C has both a name and a correlate (translation), i.e., an expression with the same meaning, in ML_C , (iii) ML_C has an undefined 1-place predicate, “ T ”, designated as the truth predicate of L_C , as well as other predicates definable as semantic predicates of L_C , and (iv) ML_C has variables of a higher-order than those of L_C (or a set-theoretical apparatus richer than that of L_C).
- (iii) *Definitions (in ML_C)*. Tarski defines the language of L_C by (mathematical) induction. His goal is a definition of “sentence” for L_C , but since simple quantifications are constructed by induction from *formulas* rather than sentences, he begins by defining the notion of a formula.

Notation: Let “ v_i ” and “ v_j ” be schematic symbols representing arbitrary variables, x_i and x_j , of L_C , and let “ Φ ”, “ Ψ ” and “ σ ” be schematic symbols representing arbitrary expressions of L_C . For each primitive constant c of L_C (logical or non-logical), let \underline{c} be a name of c in ML_C and $\underline{\underline{c}}$ an ML_C -name of the object denoted by c in ML_C (a correlate of c in ML_C or a translation of c to ML_C).

*Formula (of L_C)—Inductive Definition*¹⁰⁰

1. “ $v_i \subseteq v_j$ ” is a formula.
2. If Φ is a formula, “ $\sim \Phi$ ” is a formula.
3. If Φ and Ψ are formulas, “ $\Phi \vee \Psi$ ” is a formula.
4. If Φ is a formula, “ $\forall v_i \Phi$ ” is a formula.
5. Only expressions obtained by 1-4 are formulas.

Sentence (of L_C): σ is sentence iff σ is a formula with no free occurrences of variables (variables which are not in the scope of any quantifier).

Tarski naturally chose the recursive method to construct his definition of truth for L_C . The recursive method enables us to define predicates ranging over infinitely many objects in a finite manner, provided certain conditions are satisfied. I will not specify the conditions for recursive definitions here (for a good account see Enderton 1972/2001, Section 1.4), but the idea is that if every sentence of an inductively defined language L is uniquely generated from finitely many atomic sentences by finitely many logical operations, and if the atomic sentences and logical operators of L are finitely specifiable, then truth for L can be recursively defined. Such a definition determines the truth value of each sentence of L based on (i) the truth values of its atomic constituents, and (ii) its structure-generating operators. For example, if the only structure-generating operators of L are Negation and Disjunction, then truth for L is definable by specifying (i) the truth values of the atomic sentences of L , (ii) a rule for determining the truth value of a

¹⁰⁰ For simplicity I use regular quotation marks instead of square quotes (as in, e.g., Quine 1940/51).

negation given the truth value of the negated sentence, and (iii) a rule for determining the truth value of a disjunction given the truth values of its disjuncts.

If L contains quantifiers, however, truth for L cannot be defined in this way. Sentences involving quantifiers are, as we have pointed out above, generated not from atomic sentences but from atomic formulas, including formulas with free variables, and such formulas do not have a truth value. (For example “ $(\forall x)Px$ ” is generated from the atomic formula “ Px ”, which, having a free variable, has no truth value.) But truth for L can be recursively defined via an auxiliary notion, *satisfaction*, applicable to formulas. The notion of satisfaction is an intuitive semantic notion (in Tarski’s sense of “semantic”): The atomic formula “ x is even” is satisfied (in the domain of the natural numbers) by the objects 0, 2, 4, . . . More generally, “ Rx_1, \dots, x_n ” is satisfied by an n -tuple of objects, $\langle a_1, \dots, a_n \rangle$, iff a_1, \dots, a_n (in that order) stand in the relation R (the relation referred to by “ R ”). The definition of *truth* for L proceeds in two steps: (i) a recursive definition of *satisfaction* for L , and (ii) a (non-recursive) definition of *truth* for L based on the recursive definition of satisfaction.

Now, let g be a function in ML_C which assigns to each variable of L_C an object in the domain, D_C , of L_C . We will call g “an assignment function for L_C ” and refer to $g(v_i)$ as “ g_i ”.

Satisfaction (of a Formula of L_C by g)—Recursive Definition:

1. g satisfies “ $v_i \subseteq v_j$ ” iff $g_i \subseteq g_j$
2. g satisfies “ $\sim \Phi$ ” iff $\sim (g \text{ satisfies } \Phi)$.
3. g satisfies “ $\Phi \vee \Psi$ ” iff $[(g \text{ satisfies } \Phi) \vee (g \text{ satisfies } \Psi)]$.
4. g satisfies “ $\forall v_i \Phi$ ” iff $\forall g'$ (if g' differs from g at most in g_i , then g' satisfies Φ).

T (Truth of a Sentence of L_C): $T(\sigma)$ iff σ is a sentence and $\forall g (g \text{ satisfies } \sigma)$.¹⁰¹

Does this definition satisfy the conditions of formal and material adequacy? Tarski approaches this question in the same way he approached the definition of truth itself—through an example. To prove that his definition satisfies these conditions, Tarski says, “would require the setting up of an entirely new apparatus: in fact it involves the transition to a level one step higher—to the meta-metatheory, which would have to be preceded by the formalization of the metatheory which forms the foundation of our investigations” (Tarski 1933: 195). To avoid this complexity, he once again addresses the issue by means of an example (given in Tarski: 196).¹⁰²

One way to understand the significance of Tarski’s definition of truth is to see what consequences he draws from it and what uses he makes of it. Tarski’s next step (in his

¹⁰¹ We can replace “ $\forall g (g \text{ satisfies } \sigma)$ ” by “ $\exists g (g \text{ satisfies } \sigma)$ ”, since the two give the same results. (See standard textbooks of mathematical logic.)

¹⁰² This appears to be a recurrent choice that Tarski makes: whenever possible, simplify the discussion (though only in ways that do not involve sacrificing rigor).

1933 paper) is to study “some characteristic general theorems that can be derived from [his definition]” (Tarski: 197). It is significant that all these theorems are *metalogical*. They include “The principle of contradiction”, “The principle of excluded middle”, the principle “If $X \subseteq Tr$ [Tr being the set of true sentences] then $Cn(X) \subseteq Tr$ [$Cn(X)$ being the set of *logical* consequences of X]” (Tarski: 197–8), etc. This suggests that the center of gravity of his definition is logical.

Furthermore, Tarski’s most important use of his definition of truth is in logic. Shortly after completing his 1933 paper on truth, Tarski proceeded to develop a *general semantic definition* of the central *logical* (metalogical) notion of “*logical consequence*”, based on his general method of defining truth (Tarski 1936a). Adding the notion of “*model*” (and the derivative notion of “*truth in a model*”), his definition of logical consequence can be concisely formulated as follows:

- (LC) The sentence σ is a *logical consequence* of the set of sentences Σ iff there is no model (for the underlying language) in which all the sentences of Σ are *true* and σ is *false*.¹⁰³

Based on all of the above, it is quite clear how Tarski’s general method of defining truth for formalized language should be formulated. Limiting ourselves (for the sake of brevity) to standard first-order languages (i.e., open languages formulated within the framework of standard first-order logic), we may describe the Tarskian method of defining truth for formalized languages as follows.¹⁰⁴

Tarskian Method of Defining Truth for an Arbitrary (Standard) first-order Language, L

Let L be an arbitrary standard first-order language. We divide its vocabulary into two groups: (i) logical vocabulary, and (ii) non-logical vocabulary. The logical vocabulary includes a complete set of logical sentential connectives, identity, and the universal/existential quantifier. The non-logical vocabulary may include individual constants (c_i), m -place functions (f^m_i), and n -place predicates (P^n_i). We define “singular term”, “formula” and “sentence” of L in the usual way, assigning to each logical constant a special entry, and treating all non-logical constants of the same syntactic category under the same entry.

Turning to the semantic definition of truth for L , we proceed in four steps:

- I. *General Semantic Machinery*: U —universe of discourse of L (a non-empty set of individuals). D —a denotation function that assigns to each non-logical constant of L a denotation in (over) U of the corresponding ontological category (D assigns individuals to individual constants, m -place functions to m -place function symbols, and n -place relations to n -place predicates). G —a set of assignment functions g as in Tarski’s example.

¹⁰³ We will discuss logical consequence and models in detail in Chapter 10.

¹⁰⁴ Many details will be left out. These can be filled in by consulting standard textbooks of mathematical logic.

II. Denotation of (Singular) Terms (t) under g :¹⁰⁵

1. $g(\underline{c_i}) = D(\underline{c_i}) = \underline{c_i}$
2. $g[\underline{f^{m_i}}(t_1, \dots, t_m)] = D(\underline{f^{m_i}})[g(t_1), \dots, g(t_m)] = \underline{f^{m_i}}[g(t_1), \dots, g(t_m)]$

III. Satisfaction (of a Formula of L_C by g)—Recursive Definition: Each logical constant receives a special entry in the definition, while all non-logical constants (of the same syntactic category) are treated en block, i.e., given a single entry.¹⁰⁶

1. g satisfies " $\underline{P^n_i}(t_1, \dots, t_m)$ " iff $\langle g(t_1), \dots, g(t_m) \rangle \in D(\underline{P^n_i})$ iff $\langle g(t_1), \dots, g(t_m) \rangle \in \underline{P^n_i}$
2. g satisfies " $\underline{t_i} \equiv \underline{t_j}$ " iff $g(t_i) \equiv g(t_j)$
2. g satisfies " $\sim \Phi$ " iff $\sim (g \text{ satisfies } \Phi)$.
3. g satisfies " $\Phi \vee \Psi$ " iff $[(g \text{ satisfies } \Phi) \vee (g \text{ satisfies } \Psi)]$.
4. g satisfies " $\forall \underline{v_i} \Phi$ " iff $\underline{\forall g'}$ (if g' differs from g at most in g_i , then g' satisfies Φ).¹⁰⁷

IV. T (Truth of a Sentence of L_C)

$T(\sigma)$ iff σ is a sentence and $\underline{\forall g}(g \text{ satisfies } \sigma)$.

THE PHILOSOPHICAL SIGNIFICANCE OF TARSKI'S DEFINITION

While Tarski's theory of truth is widely viewed as one of the prime achievements of twentieth-century analytic philosophy, its philosophical significance has been repeatedly questioned.¹⁰⁸ My own suggestion is that philosophically, Tarski's theory should be viewed as an account of the *contribution of logical structure to truth*. Tarski's theory tells us how the logical structure of a given sentence affects its truth value, not how other types of structure/content (modal, physical...) do. Tarski's theory is a theory of a specific, albeit basic and general constituent of truth, namely, its logical constituent. Its goal is to describe, in an exhaustive, systematic, and informative manner, that part of the truth conditions of sentences which is due to their logical structure. This explains why Tarski's theory of *truth* is so important and fruitful in *logic*. Its importance to logic is one thing that most philosophers agree about: lying at the heart of the semantic branch of logic, Tarski's theory makes an enormous contribution to logic.

Tarski's theory does not make similar contributions to other disciplines. While Tarski's definition of truth for a language L yields, almost all by itself, a definition of *logical consequence* for L (assuming ML has a sufficiently rich set-theoretical

¹⁰⁵ t is an individual constant or a variable or a functional expression of the form $f^m(t_1, \dots, t_m)$.

¹⁰⁶ In languages with a number of logical constants of the same type, each is usually given its own entry. Although technically it is possible to treat all logical constants of the same syntactic type together, treating them separately mirrors their special standing in the Tarskian definition. We will discuss the special status of logical constants in Tarski's definition in more detail shortly.

¹⁰⁷ Note: in 1–4, a double underline indicates the ML version of an L -Symbol.

¹⁰⁸ Note that for the sake of simplicity, we will no longer explicitly distinguish between a symbol, S , of L , its name in ML , and its correlate (expression with the same meaning) in ML , or between Tarski's definition of truth and his method of defining truth, unless these are essential for the discussion.

apparatus), it does not yield definitions of *epistemic*, *modal*, *physical*, or *biological consequence* for L. That is, it does not lie at the heart of consequences like:

- (7) *a* knows that P; therefore: *a* believes that P.
- (8) Necessarily P; therefore: Possibly P.
- (9) The force exerted on body *a* at time *t* is zero; therefore: The acceleration of *a* at *t* is zero.
- (10) *a* is a human female; therefore: *a* does not have a Y chromosome.¹⁰⁹

Why does Tarski's theory yield an account of *logical* consequence but not an account of other types of consequence? What features should a theory of truth have in order to yield a concept of consequence of type X?

Consequence relations are relations that guarantee preservation (transmission) of truth, relations that guarantee that if the premises are all true, the conclusion is true too. How do they do this? Leaving the *philosophical grounding* of logical and other types of consequence to Chapter 10, we can say that on the *linguistic* level consequence relations are associated with a set of *distinguished* terms. The distinguished terms in (7) are "knows" and "believes", in (8)—"necessarily" and "possibly", in (9)—"force", "time", "acceleration", and in (10)—"human", "female", and "Y-chromosome". These terms are *held fixed* in (7)–(10), and in this way serve as the basis for the inferences drawn. (For an extended discussion see Sher 1996.) In contrast, the distinguished terms of *logical* consequences are *logical* constants, and it is these constants, and *only* these constants, that have a *distinguished* status in Tarski's theory.

The distinguished status of logical constants is often taken as an *obvious* feature of formalized, and even natural, languages—something that can be taken for granted and requires no explanation. But in principle there are many ways to regiment and formalize a given discourse. One can assign a privileged status to modal, epistemic, physical, or biological constants, using those to generate the dominant structure or form of the sentences involved. In the Tarskian definition it is the *logical* constants that are assigned this role. And it is due to this fact that the Tarskian definition of *truth simpliciter* yields a definition of *logical consequence (and truth)*. The Tarskian definition of logical consequence (and other logical notions) incorporates a definition of truth as a major component. But not any definition of truth would do. A definition that assigns the role of distinguished terms to a special category of non-logical constants will not. Had the distinguished constants of the Tarskian definition been epistemic or biological, it would have given rise to the notion of epistemic or biological consequence; as it is, it gives rise to the notion of *logical* consequence.

¹⁰⁹ Note that (9) and (10) are viewed as complete inferences, not as enthymemes. (For a similar approach, see Brandom 1994a.)

Tarski's theory of truth thus provides an account of a factor of truth that is both *specific* and *general*. A vivid example of both the pervasiveness and the partial nature of the logical factor in truth is the following passage from Darwin:

He who believes in the struggle for existence and in the principle of natural selection, will acknowledge that every organic being is constantly endeavouring to increase in numbers; and that if any one being varies ever so little, either in habits or structure, and thus gains an advantage over some other inhabitant of the same country, it will seize on the place of that inhabitant, however different that may be from its own place. Hence it will cause him no surprise that there should be geese and frigate-birds with webbed feet, living on the dry land and rarely alighting on the water, that there should be long-toed corncrakes, living in meadows instead of in swamps; that there should be woodpeckers where hardly a tree grows; that there should be diving thrushes and diving Hymenoptera, and petrels with the habits of auks (Darwin 1859: 227).

Notice how prevalent the logical factor is in this passage, yet how partial its influence is. Its linguistic constituents, logical expressions, are indispensable for making Darwin's point, yet they are intertwined with a host of other expressions which are just as important for the content of this passage as the logical expressions. The logical factor is interlocked with other factors, and the *truth* (*falsity*, general validity or invalidity) of the cited paragraph is the combined result of all these factors.¹¹⁰

The clue to identifying the distinguished constants of Tarski's definition is their *fixity*¹¹¹ and *specificity*. While the non-logical constants vary from one formalized (Tarskian) language to another, the logical constants remain the same. This is the reason that *general renditions* of Tarski's definition represent the non-logical constants by schematic symbols. The schematic symbols represent constants which vary from one language to another (within the range of Tarski's method), while the non-schematic symbols represent constants which are fixed across languages. Each fixed constant has a precise, individual satisfaction condition, while the non-logical constants are assigned satisfaction conditions in bulk. We may say that the *fixed* or *distinguished* elements of Tarski's definition are those elements which are *singled out* in the *syntax* for a *fixed, specific treatment* in the *semantics*. The non-distinguished constants, in contrast, are treated "en masse": all constants of the same grammatical category are grouped under a single schematic entry, and in different languages this entry is instantiated by different denotation/satisfaction conditions. The role of non-distinguished constants in Tarski's definition is thus auxiliary. They are limited to the role of arguments of *logical* operators, and their schematic status enables Tarski's method to apply to a *diverse range of languages* (with the same logical apparatus but different non-logical vocabulary).

¹¹⁰ The exact demarcation of the logical constants in this passage is not important; further analysis might lead us to change it, but the principle will not change.

¹¹¹ This is the same kind of *fixity* that we talked about in Chapter 3, Sections 4–5, and Chapter 4, Sections 1–2.

Now, in offering this interpretation of Tarski's theory, one has to distinguish two separate goals: (a) an exegetical goal, (b) a constructive goal. Our main goal in this essay is constructive, and although I believe that the present construal is exegetically significant as well, it is primarily designed to point to a central component of the *logicality principle* of truth (as it is conceived in our theory).

II. Immunity to criticisms of Tarski's theory

Tarski's theory has been subjected to multiple criticisms, four of which we will address here. These can be divided into two groups: (a) criticisms claiming that Tarski's definition of truth is (i) relative to language and (ii) trivial; and (b) criticisms objecting to (i) its restriction to formalized languages and (ii) its hierarchical treatment of truth. In addressing these criticisms, my goal is to show that these criticisms do not affect our own theory, including our own construal of Tarski's theory. If, and to the extent that this construal reflects Tarski's own understanding of his theory, it is also a defense of Tarski himself, at least against the (a) criticisms. (Our response to the (b) criticisms is based on our own overall approach to truth and knowledge, which is independent of Tarski.)

RELATIVITY TO LANGUAGE AND TRIVIALITY

The relativity-to-language criticism (e.g., Blackburn 1984) says that whereas a general theory or definition of truth should hold for all languages, Tarski's theory/definition of truth does not. There is a Tarskian definition of truth for language L_1 , a different Tarskian definition of truth for language L_2 , etc. The triviality criticism (Field 1972) says that the *atomic* entries in Tarski's definition (the entries for expressions with no logical constants) are trivial and that this renders his entire definition trivial.

Both the relativity-to-language criticism and the triviality criticism, however, are directed at a construal of Tarski's theory that is very different from ours. On our construal, Tarski's theory is a *partial* theory of truth, a theory of a *special* aspect of truth—the contribution of *logical structure* to the truth value of sentences. The burden of *generality* (*unity*), and *informativeness* (*non-triviality*) therefore falls on the *logical* entries of the definition, i.e., the entries for expressions governed by a logical constant (identity formulas, negations, disjunctions, and universal quantifications). Since the logical constants are the same in all Tarskian languages and the entries assigned to them in Tarski's definition are the same in all languages, Tarski's definition is *not* (or not significantly) relative to language on our construal. Similarly, since the burden of informativeness in Tarski's definition, as construed by us, falls on the logical entries, the triviality of the non-logical entries is not sufficient to trivialize the definition as a whole. To the extent that Tarski's definition is a definition of the contribution of *logical* structure to truth, the role of the *non-logical* entries is merely *auxiliary*. As a result, no more than a "surface" (and possibly trivial) account of non-logical reference, satisfaction, and truth is needed.

In contrast, the relativity and triviality criticisms do hold for other construals of Tarski's theory. In particular, they hold for the prevalent construal of his theory as a *general, reductionist* theory of truth. On this construal, Tarski's theory is, or purports to be, a *general, exhaustive* theory of truth, and its main achievement lies in its *reduction* of the truth conditions of *all* sentences of a given language to the satisfaction conditions of a relatively small collection of expressions of this language, namely, its *atomic* formulas (in an unproblematic way)¹¹². Under this construal, the burden of generality and informativeness falls on the *atomic (non-logical)* entries of Tarski's definition, namely, the reference and satisfaction entries for *non-logical* atomic expressions. Since these change from one Tarskian language to another, the "generalist" version of Tarski's definition is (significantly) relative to language. The same holds for the triviality criticism. Since the burden of substantiveness in the generalist-reductionist version of Tarski's definition falls on the non-logical atomic entries, and since these entries are indeed uninformative, the definition as a whole is trivial or uninformative.¹¹³

Neither criticism, as we have explained earlier, holds for our construal of Tarski's theory. Indeed, our discussion of the methodological difficulties facing the theory of truth (Chapter 7, Section 2) explains why the attempt to render *all* the atomic entries of a general definition of truth informative is pointless. The enormous diversity of truth (truth conditions) is almost completely concentrated in the atomic domain, and therefore an informative definition of truth for this (entire) domain is highly problematic. Abstract and concrete sentences, mathematical, physical, moral, and a host of other sentences populate this domain; as a result, the tension between generality and specificity is at its peak in it, and no more than a schematic, non-substantive treatment of its sentences is feasible. In contrast, an informative account of a single factor of truth (or of a small cluster of homogeneous factors) is possible, and this is what Tarski's definition is designed to accomplish on our construal. Tarski's theory is designed to account for the *logical factor* in truth in a general and informative manner, leaving the informative treatment of other factors to other branches of the theory of truth. The informativeness of Tarski's definition therefore depends on its logical entries.¹¹⁴

¹¹² "Unproblematic" varies from critic to critic. For Field (1972) an unproblematic way accords with the principles of *physicalism*; others set weaker requirements.

¹¹³ Notice that in general a recursive definition allows two interpretations: an interpretation that regards the atomic entries as its center of gravity, and interpretation that regards the recursive entries as the center of gravity. But it also allows a combination of the two, as in the case of Tarski's definition when the logical constants include the identity relation. In that case, the non-recursive entry for identity and the recursive entries for the other logical constants are the main entries.

¹¹⁴ Whether the logical entries in Tarski's definition, as they are commonly formulated, are as informative as they should be is an open question. (See discussion of "The logicity principle beyond Tarski" later on in this section.)

FORMALIZED LANGUAGE AND HIERARCHICAL TREATMENT OF TRUTH

The formalized-language and hierarchical-treatment-of-truth criticisms concern (i) Tarski's restriction of his truth method to *open* and formalized languages¹¹⁵, and (ii) the hierarchical nature of his method: the fact that the truth predicate of a language *L* belongs to its metalanguage, *ML*, the truth predicate of *ML* belongs to *MML*, and so on, and the related fact that the definition of truth for *L* is given in *ML*, the one for *ML* in *MML*, and so on. Tarski justified these features of his method as effective devices for blocking the Liar paradox (and other semantic paradoxes). But many philosophers claim that these devices are both unacceptable and unnecessary. They are unacceptable because they are too restrictive, too ad hoc, and subject to too many serious problems; and they are unnecessary because it is possible to define truth for *closed*, natural, languages in a non-hierarchical manner without generating a paradox. Let me elaborate.

Concerning (i), philosophers object on the ground that our most important language is natural language, not open formalized languages, and that an adequate definition of truth must be applicable to this language. They are especially adamant that in natural language there is only one truth predicate, "true", not a series of distinct predicates—"true₁", "true₂", "true₃", ...,—and they insist that an adequate theory of truth must preserve this feature. Concerning (ii), philosophers say that Tarski's hierarchy is ad hoc, that "Tarski has no independent reason for postulating the distinction between object language and metalanguage other than to solve the Liar Paradox" (Kirkham 1992: 281). They further say that Tarski's solution to the Liar paradox is too strict: it eliminates not only paradoxical uses of "true" and other semantic notions, but also innocuous uses of these notions. For example, Tarski's theory is incapable, in principle, of dealing with sentences whose level in the Tarskian hierarchy is not fixed in advance. A case in question, suggested by Kripke, is the following: Consider the two sentences:

- (11) All of Nixon's utterances about Watergate are false (Kripke 1975: 695)

and

- (12) Everything Dean says about Watergate is false (Kripke: 696)

where (11) is (presumably) uttered by Dean and (12) by Nixon during the Watergate crisis, and where each includes the other in the range of its initial quantifier. This pair of sentences is unproblematic in the sense that there is an intuitive model for it, i.e., a reasonable scenario in which both sentences are assigned a definite truth value without collapsing into paradox. Yet there is no place for this pair in Tarski's hierarchy: (11) must be higher up in the hierarchy than (12), and vice versa.

¹¹⁵ Reminder: Open languages are languages that do not contain names of their own expressions or semantic predicates applicable to their own expressions.

Philosophers also say that Tarski's definition implicates truth in an infinite regress: "true₁" is defined in terms of "true₂", "true₂" is defined in terms of "true₃", and so on. And they say that Tarski's hierarchical account of truth sets unreasonable limits on our theoretical goals. An especially broad outlook on the matter is offered by McGee:

At issue is the possibility of a unified scientific understanding in which human thought and action are no less intelligible or more mysterious than the planetary orbits. If we adopt [Tarski's] proposed solution, we shall find that within the object language we are unable even to describe human thought and action. . . . [I]ntentional human activities, such as speaking, believing, . . . will be indescribable and inexplicable. Within the metalanguage we can obtain fragmentary descriptions of human thought and actions. . . . [B]ut thought about thought and talk about talk will remain indescribable and inexplicable. Thus, if we accept the limitations imposed by Tarski's proposal for avoiding antinomies, we forfeit one of the highest aspirations of the human spirit, the aspiration to self-understanding (McGee 1991: 79).

Philosophers further point to the existence of multiple solutions to the Liar paradox that do not involve a hierarchy of languages (see Kripke 1975, McGee 1991, Gupta and Belnap 1993, among others). But a note accompanying one of these solutions directs our attention to a remarkable fact:

It seems likely that many who have worked on the [non-hierarchical] approach to the semantic paradoxes have hoped for a universal language, one in which everything that can be stated at all can be expressed. . . . Now the languages of the present approach contain their own truth predicates and even their own satisfaction predicates, and thus to this extent the hope has been realized. Nevertheless the present approach certainly does not claim to give a universal language, and **I doubt that such a goal can be achieved**. First, the induction defining the minimal fixed point is carried out in a set-theoretic **meta-language**, not in the object language itself. Second, **there are assertions** we can make about the object language **which we cannot make in the object language**. For example, Liar sentences are *not true* in the object language, in the sense that the inductive process never makes them true; but we are precluded from saying this in the object language by our interpretation of negation and the truth predicate. If we think of the minimal fixed point, say under the Kleene valuation, as giving a model of natural language, then the sense in which we can say, in natural language, that a Liar sentence is not true must be thought of as associated with some **later stage** in the development of natural language, one in which speakers reflect on the generation process leading to the minimal fixed point. It is not itself a part of that process. The **necessity to ascend to a metalanguage** may be one of the weaknesses of the present theory. **The ghost of the Tarski hierarchy is still with us** (Kripke 1975: 714).

Kripke is right about the deep entrenchment of Tarski's hierarchy (or something like it) in human thought and theorizing. But from our perspective there is nothing ghostly or undesirable about this. On the contrary. Tarski's conception of truth as involving transcendence to a metalanguage is important beyond the problem of the paradoxes or even the logicity principle. The Tarskian hierarchy reflects the

dynamic nature of human thought and knowledge; in particular, our ability to move up and down, back and forth, from one language to another, in cognizing the world (including ourselves in it). This is the insight behind Tarski's hierarchy.

Focusing on the dynamics of the Tarskian hierarchy, we see the three components of the fundamental principle of truth—*immanence*, *transcendence*, and *normativity*—at work. First, every language *L* in the Tarskian hierarchy is an *immanent* language, i.e., a language whose sentences say something about the world, or some aspects of the world (languages, too, being such aspects). Second, to construct a definition of truth for *L*, we *transcend* *L* to the next language in the hierarchy, *ML*, a language from which we can pose, and in principle answer, the *normative question of truth* as it applies to sentences of *L*. Since the question of truth is essentially a *correspondence* question, i.e., a question about whether a given sentence, *S*, of *L* measures up to reality, *ML* is required to have in view both *L* and those aspects of the world that *L* can relate to. This is accomplished by endowing *ML* with both resources for forming names of all expressions of *L* (including all sentences of *L*) and resources for talking about the world, including all those aspects of the world that *L* can talk about. Finally, in order to state the conditions under which a given sentence of *L* is true in accordance with the *correspondence* principle, *ML* has to have resources for saying that a given sentence of *L* is true *iff* it corresponds to reality, i.e., *iff* certain conditions concerning its relation to the world and the way the world is are satisfied. And indeed, Tarski designs *ML* in a way that ensures it has all these resources. (For a full paper analysis and defense of Tarski's hierarchical solution to the Liar paradox from the perspective of the present essay see my "Truth and Transcendence: Turning the Tables on the Liar Paradox" (2016a).)

Turning to Kripke's example of sentences that cannot find their proper place in the Tarskian hierarchy, what this example shows, from our perspective, is that for some purposes the Tarskian hierarchy, as it now stands, is *too static*. What we need is a *more dynamic* hierarchy, one in which non-paradoxical sentences can eventually find their place in the hierarchy, though not necessarily in one step or at the outset.

As for a definition of truth for *closed* languages, including *natural language*, it is instructive that most attempts to solve the Liar paradox within the framework of natural language involve an *internal hierarchy*, or a *dynamic process*, or both (as in Kripke's solution). From our perspective, such solutions are perfectly acceptable.¹¹⁶ In this essay, however, we are primarily interested in truth from an epistemic perspective, so the importance of natural language is secondary for us.

Concerning the general theoretical adequacy that McGee justly requires, we can satisfy this requirement by invoking the holistic methodology, dynamic model of knowledge, transcendence, and other tools we have developed in this essay. These

¹¹⁶ The non-bivalence often involved in such solutions does not pose a problem for us, since as we have already noted (and will briefly note again in Chapter 10, Section 4) our conception of truth, knowledge, and logic is sufficiently broad and flexible to accommodate non-bivalence (if and when this is appropriate).

enable us to obtain a unified perspective on human thought and action without adopting a static, one-dimensional perspective on it.

The same tools enable us to avoid the infinite-regress objection (noted earlier). In Chapter 2, Section 2 we have shown how, working within a Tarskian hierarchy, we can step aside and, having the entire hierarchy in view, understand its general principle. That is, once we figure out how to define truth for one language in the hierarchy by ascending to the next language in the hierarchy, we know how to do so for any language in the hierarchy, regardless of its level. Since the principle remains the same as we ascend to higher levels (the same detailed entries for the logical vocabulary and the same schematic entries for the non-logical vocabulary), we do not need to complete the hierarchy in order to master Tarski's method.¹¹⁷ Of course, this realization is arrived at from some standpoint within some Tarskian hierarchy. But by continuing to move sideways to yet other Tarskian hierarchies we quickly realize that any two Tarskian hierarchies are essentially the same. So, by learning how to master the method of defining truth in one Tarskian hierarchy, we have learned how to master it for any Tarskian hierarchy. Of course, we must be cautious in drawing this lesson: we must subject it to critical examination and make needed adjustments as we go. But infinite regress by itself need not be an insurmountable obstacle to knowledge, given holism—not if we use appropriate means to deflect it.¹¹⁸

III. The logicity principle beyond Tarski

While Tarski's (1933) theory makes an invaluable contribution to our understanding of the logicity principle of truth, it leaves much philosophical work to be done. (a) We need an informative formulation of the satisfaction conditions of the logical constants and the formulas they govern (and not just a metalinguistic transliteration of these expressions as in Tarski's original definition). (b) We need to understand the nature of the logical constants. This includes a description of their distinctive characteristics, a criterion distinguishing logical from non-logical constants, and a philosophical ground for this criterion.¹¹⁹ (c) We need to understand how the truth conditions of logically complex sentences fit in with a general correspondence account of truth. Are the logical constants referring constants, and if so, what properties/relations do they refer to? Are the satisfaction conditions of logically complex formulas correspondence conditions, and if so, what kind of correspondence conditions are they? What are the truth

¹¹⁷ This situation is analogous in certain respects to that of being able to attain a fairly comprehensive understanding of the natural numbers without introducing infinity.

¹¹⁸ For further discussion of some of the issues raised here see Sher (2016a).

¹¹⁹ In his 1933 paper Tarski did not offer a systematic account of the logical operators. In his 1936a paper he realized that such an account was essential for his definition of logical consequence (and ipso facto logical truth). And in his 1966 lecture (first published in 1986) he offered a criterion for logical operators, though without a philosophical rationale. (See discussions in Chapter 10.)

conditions of logical truths? What is the route of correspondence of such truths? And so on.¹²⁰

The present theory provides a universal standard of truth for the model of knowledge developed in Part II of this essay. This standard is both robust and flexible: its robustness is manifested in the requirement that true statements/theories be connected to reality in a substantial and systematic manner; its flexibility is reflected in the fact that it allows multiple variations in the form these connections take. The standard takes into account both differences in aspects of the world that a given theory is concerned with and differences in our cognitive resources for accessing different aspects of reality. This means that the standard is not a traditional correspondence standard, but a standard of “manifold correspondence”, one that allows multiple “routes” or “forms” of correspondence, including “composite” (complex, intricate) routes. In accordance with our overall approach to philosophical theorizing, the theory of truth leaves it an open question how correspondence works in any given field of knowledge and how correspondence in one field interacts with that in another. These things are to be discovered based on specific investigations rather than decided in advance.

In the next part we will use some of the theoretical tools developed earlier in this essay to take up a philosophical problem whose importance is widely recognized but whose solution is often considered unattainable, namely, the *foundational problem of logic*. The task of providing a systematic philosophical foundation for logic, and in particular a foundation concerned with the *veridicality* of logic, is one that most contemporary philosophers have given up on. This is not to say that contemporary philosophers have lost interest in understanding logic; on the contrary, their considerable interest is evidenced by current work in the philosophy of logic (see References). But rarely do contemporary philosophers confront the question of *truth* as it applies to logical theories. Suppose someone comes up with a new system of inference and suggests that we use it as our logical theory. How do we determine whether this theory is true or false? For example, how do we determine whether inferences licensed by this theory *in fact* transmit truth from premises to conclusion with a strong modal force? A veridical foundation for logic must be able tell us how to go about answering this question. It must tell us what the veridicality of a given logical theory is grounded in. It must explain why not just any “logical system” whatsoever, not just any “inferential” game that someone might invent, would work

¹²⁰ In *The Bounds of Logic* (1991), Chapter 4, I provide an informative formulation of the satisfaction conditions of the logical constants and the formulas they govern, as required by (a); in Chapter 3 of the same book I provide a descriptive characterization and a criterion of logical constants; and in Chapter 10 of this essay I will answer the questions posed by (c).

in the world. It must explain the difference between a system of logical rules that satisfies the veridicality requirement and one that does not.

The epistemic theory developed so far in this essay provides us with resources for doing all this, and the resulting foundation will, in turn, further support and expand the theory. It will explain why logic is normatively bound to lie not just in the center but also—and significantly so—in the periphery of our system of knowledge, what its interface with reality is, how it meets the universal friction requirements, and in what way logical truth is based on correspondence. And it will further provide an account of logic's role in knowledge, the relation between logic and mathematics, the normativity of logic, the characteristic traits of logic (topic neutrality, generality, strong modal force, etc.), error and revision in logic, the scope of logic, pluralism with respect to logic, and so on.

PART IV

An Outline of a Foundation for Logic

9

The Foundational Problem of Logic

9.1 The Foundational Problem of Logic as a Methodological Problem

It is an interesting and a puzzling fact that a systematic philosophical foundation for logic has rarely been attempted.¹ What I have in mind is a unified theoretical foundation, focused on logic itself rather than on mathematics, science, or natural language. Such a foundation would be critical and explanatory, and it would be especially concerned with the *veridicality* of logic. In connection with this concern it would critically examine the basic features of logic, the tasks logic is designed to perform, the source of truth and falsehood of logical and metalogical claims,² the grounds on which logical theories should be accepted (rejected, or revised), the ways logical theories are both constrained and enabled by the human mind on the one hand and the world on the other, the relations between logic and related disciplines (e.g., mathematics), the normativity of logic, and so on. The list is, in principle, open-ended since new interests and concerns may be raised by different persons and communities at present and in the future, but veridicality is at the core of such a foundation.

Now, the surprising fact is not that philosophers are not interested in the foundation of logic. They are. Many philosophers have strong views, opinions, feelings, intuitions, etc., about many foundational issues concerning logic. Furthermore, many contemporary philosophers have written penetratingly on such questions as whether logic is *apriori* or empirical, whether there is one logic or many, whether logic is bivalent or not, the nature of the law of non-contradiction, the nature of the formality of logic, the normativity of logic, logic and knowledge, and so on. Still, with a few exceptions—e.g., Hanna (2006) and Maddy (2007, Part III)—a systematic and unified foundation for logic that deals with all or most of these issues and, most

¹ A similar sentiment was recently expressed by Maddy (2012: 482–3).

² From specific object-language claims such as “No object is both round and not round” to general metalinguistic claims, for example, the claim that a certain inference-form is logically valid.

importantly, with the *veridicality* of logic, has rarely been attempted (at least within the analytic stream of philosophy).³

A foundational theory of the kind we are seeking does not purport to be infallible. Like all other human theories it is subject to standards appropriate to its field (in the present case, philosophy and metalogic) and open to criticisms and improvements. The foundation it develops is a foundation for logic in a broad sense—the discipline of logic rather than a specific logical theory—though such a foundation would desirably provide us with tools for criticizing, justifying, evaluating, constructing, and improving specific logical theories (or proposals for such theories) as well.

My motivation in pursuing such a foundational project is both general and particular. On the one hand, my interest in the foundation of logic is part of my interest in the foundation of knowledge in general. On the other hand, my interest is specific to logic, due to its unique features and its special standing in our system of knowledge. This special standing has been widely recognized. Compare logic and physics, for example. It is widely recognized that physics is bound by the laws of logic, but logic is (at least for the most part) not bound by the laws of physics. Logic is conceivable outside physics, but physics is inconceivable outside logic. Logic abstracts from the content of physical terms, but physics does not abstract from the content of logical terms. Logic provides tools for physical theories, but physics (for the most part) does not provide tools for logical theories. A serious error in logic is likely to undermine physics, but serious errors in physics are unlikely to undermine logic. Logic delineates some of the most basic forms of human thought and its expression, provides us with the most basic tools of valid inference, tells us what theories are inconsistent, etc. Logical form, logical inference, logical criteria of consistency, are all ingredients that no system of knowledge can do without. Our system of knowledge can survive the removal of many a science, but not of logic. As a result, a foundation or a grounding for logic is crucial for our system.⁴

Given the many issues a foundational theory of logic must address, however, there is no question of encompassing their full range in this essay, let alone in a single part of it. What I am looking for, instead, is a fruitful standpoint from which to approach this project and a constructive investigation that addresses some of its key questions in a unified manner. Such an investigation would serve as a starting point for a more complete foundation and, hopefully, as a catalyst for further theoretical discussion of the foundation of logic within the philosophical and logical communities. Furthermore, logic, today, has multiple branches, and the present investigation cannot be directed at all of them. Rather, it focuses on one central branch of logic, a branch that today is often referred to as “mathematical logic” and in earlier times took the forms

³ There are several points of similarity between the present work and Hanna’s and Maddy’s works. But these works differ from it in being radically psychological/naturalistic. (I discussed Hanna’s work in Sher 2007.)

⁴ In this Part, grounding will be a recurrent theme. See “A Note on Grounding” (Chapter 2, Section 2).

of syllogistic logic, Fregean logic, and type-theoretic logic. It will, nevertheless, address the plurality of logics, and it can be viewed as an example of how a foundational investigation of any branch of logic might proceed.

I have said that systematic attempts to construct a philosophical foundation for logic have been rare. But was not the period between the late nineteenth century and the early twentieth century a period of “foundational studies in logic and mathematics”, indeed a period of extraordinary growth and remarkable breakthroughs in this area? The answer is “Yes” with a caveat. Yes, there were foundational investigations and groundbreaking developments, but for the most part they aimed at a foundation for *mathematics*, with logic playing a largely instrumental, if crucial, role. Frege, for example, developed a logical system that would provide a foundation for mathematics, but aside from a few hints, did not attempt to provide a systematic philosophical foundation for logic itself.⁵ Russell improved and further developed Frege’s logicism, but although he appreciated the need to provide a systematic philosophical explanation of logic itself—one that would answer such questions as: “In virtue of what are logical propositions true?”—he despaired of accomplishing this task. Thus he says:

The fundamental characteristic of logic, obviously, is that which is indicated when we say that logical propositions are true in virtue of their form. . . . I confess, however, that I am unable to give any clear account of what is meant by saying that a proposition is “true in virtue of its form” (Russell 1938: xii).

Indeed, many of the momentous discoveries in metalogic (by Hilbert, Gödel, Turing, and others) are commonly designated as contributions to “metamathematics”. These epochal achievements, however, are not irrelevant to the foundational problem of logic. On the contrary, by giving rise to a sophisticated logical framework and establishing its mathematical properties, they created a fertile ground for addressing this problem.

It is all the more surprising, therefore, that relatively few philosophers have taken on this project. Most contemporary philosophers seem to believe that a substantive, theoretical foundation for logic, and especially one focused on its veridicality, is impossible. Some consider it superfluous. Quite a few are content to simply say that logic is obvious. Others view logic as conventional, hence not in need of a foundation. And so on.

This tendency to disengage from the philosophical foundation of logic goes back to the great foundational systems of philosophy. Take Kant, for example.⁶ Without offering a scholarly exegesis of Kant’s philosophy of logic,⁷ we may note that Kant’s approach to logic is quite different from his approach to other disciplines. While

⁵ For a nuanced account of Frege’s approach to the foundation of logic, see Burge (1998).

⁶ In speaking of Kant’s view of logic I focus on his view of the counterpart of current mathematical logic in his time, or what he characterizes as general, pure, formal logic (see Kant 1781/7: A52–4/B76–8).

⁷ For such exegeses, see e.g., Longueness (1993/8) and Tolley (2006 and 2013).

Kant set out to provide a substantive foundation for scientific and mathematical knowledge, he took formal logic largely as given. Kant's "Copernican Revolution" did not include a revolution in logic, and Kant saw no need for a revolution in that field. Indeed, he saw no need for a critical justification of the basic concepts and laws of logic. Logic, Kant seems to have held, participates in the veridical foundation of all other disciplines, but logic itself does not require a veridical foundation. The fact that logic had not undergone a major revision since Aristotle was viewed by Kant as highly significant: logic did not require a new foundation. Kant did see a need for further clarifications and adjustments in logic,⁸ but he did not worry about its certainty (in the way he worried about the certainty of science, for example). Thus he says:

That logic has already, from the earliest times, proceeded upon [a] sure path is evidenced by the fact that since Aristotle it has not required to retrace a single step, unless, indeed, we care to count as improvements the removal of certain needless subtleties or the clearer exposition of its recognised teaching, features which concern the elegance rather than the certainty of the science (Kant 1781/7: Bviii).

Kant, who asked penetrating questions about science, did not ask similar questions about logic. He did not ask whether any logic whatsoever, if built into a human mind, would succeed in working in the world, whether any system of rules treated as "logic" could serve as a basis for scientific knowledge, whether there are serious epistemic constraints on our logical theory, what forms of "resistance" ("friction", in our terminology) our logical theory has to meet, and so on.⁹

The scarcity of attempts to provide a veridical foundation for logic, both now and in the past, is especially surprising in light of the wide agreement among philosophers about logic's distinguished place in knowledge. If physics is dependent on logic more than logic is on physics, if a serious error in logic might undermine physics while a serious error in physics would not (or is far less likely to) undermine logic, then a foundation for logic is more crucial for knowledge than a foundation for physics. The more general, basic, and normative a given field of knowledge is, the more important it is to establish its veridicality, develop methods for identifying and correcting its errors, explain its relations to other fields of knowledge, determine its scope, etc.—in short, provide it with a foundation. But in spite of the fact that no discipline is (or at least appears to be) more fundamental than logic, a systematic theoretical foundation for logic, and in particular a non-trivial veridical foundation, has rarely been attempted. Why?

It is clear that the failure to construct a substantive foundation for logic is not due to neglect, oversight, or intellectual limitations. The extraordinary advances in logic and metalogic, the wealth of attempts to construct a philosophical foundation for mathematics and science, and the rich literature on logic itself suggest that neither

⁸ Some of which are discussed in his *Lectures on Logic* (1770s–1800).

⁹ This further elaborates a point we briefly noted in Chapter 1, Section 1.

neglect nor intellectual handicaps are the problem. Still, the absence of a veridical foundation for logic is, to use Wright's words, a "scandal". Looking back to Kant, Wright says:¹⁰

In a famous footnote, Kant wrote: 'It still remains a scandal to philosophy . . . that the existence of things outside us . . . must be accepted merely on *faith*, and that, if anyone thinks good to doubt their existence, we are unable to counter his doubts by any satisfactory proof'.

It is, if anything, a yet greater 'scandal' that we have so far acquired so little understanding of the basic epistemological architecture of logical inference (Wright 2001: 81).

What is the source of the problem? Why have there not been more systematic attempts to construct a theoretical foundation for logic? In my view, the source of the problem is *methodological*. Certain features of the customary foundational methodology make it very problematic to construct a veridical foundation for logic, and the first step in confronting the foundational problem of logic is, therefore, to overcome this methodological obstacle.

9.2 Analysis of the Problem and Alleged Remedies

I. *Foundation vs. foundationalism*

In Chapter 2 we identified a serious obstacle to a veridical foundation for knowledge in general, namely, the confounding of "foundation" with "foundationalism". This obstacle is especially prohibitive in the case of "basic" disciplines, i.e., disciplines lying at the base of the foundationalist tree or pyramid. In the case of logic, the assimilation of the two notions has the effect of turning the foundational problem into an *either-or* problem: *either* logic can be given a foundationalist grounding—which it cannot—or it must be given a "free pass". At best, logic can be given a pragmatist foundation, one that neither intends nor is able to meet the challenge of *veridicality*. One expression of this attitude is the prevalent view that a foundation for logic must be certain. Another is the deep-seated view that a foundation cannot involve any measure of circularity. The crux of the matter, here, is the so-called "logocentric" predicament": "*In order to give an account of logic, we must presuppose and employ logic*" (Sheffer 1926: 228). In Wittgenstein's idiom: to provide a foundation for logic we have to stand "somewhere outside logic", *but there is no cognitive standpoint outside logic* (Wittgenstein 1921: 4.12). Indeed, even those who are willing to accept innocuous circularity in other contexts regard the foundational project of logic as undermined by circularity. Thus, van Benthem, who is *not* a foundationalist, nevertheless regards circularity as a major problem for "*an independent foundational account of logicality*" (van Benthem 2002: 429). Like many other non-foundationalists, he seems to identify the foundational project with the foundationalist project, believing a

¹⁰ For **boldface** within citations, see fn. 3, Chapter 1.

foundational account that involves circularity is self-undermining. Furthermore, even those who do allow circularity in a justificatory account of logic seem to view such an account as falling short of a full-fledged foundation. For example, Boghossian says, with respect to his own justificatory account of logic, which admits circularity, that the justification it offers falls short of a *genuine foundational* justification, one that can “silence skeptical doubts” (Boghossian 2000: 254). He concludes his discussion with the concession that “with respect to something as **basic** as logic, [foundational justification] was **never in prospect** anyway” (Boghossian: 254).

Now, the foundational-holistic methodology developed in Chapter 2, Section 2 enables us to leave behind all these remnants of the foundationalist approach. The idea that a veridical foundation of logic must be absolutely certain is relinquished. Similarly, the total ban on circularity is relaxed. One of the main problems with the foundationalist approach, on our analysis, is its strict ordering of all fields of knowledge. The imposition of a strict ordering gives rise to the *basic-knowledge predicament*, which rules out a veridical foundation for any basic discipline, including logic. The point is that since no unit of knowledge lies below logic in the foundationalist hierarchy, no unit (or combination of units) has, or can produce, resources for grounding logic. In Chapter 2, Section 1 we presented this predicament in an inference form that would apply to logic as follows:

1. It is a central principle of the foundationalist methodology that to ground X we must use resources more basic than those generated by X.
2. It is also a central principle of the foundationalist methodology that our system of knowledge cannot generate more basic resources than those generated by the basic units.
3. Logic is a basic discipline.

Conclusion: No unit of knowledge can produce resources for grounding logic.

Logic, the foundationalist says, can take part in providing a foundation for other branches of knowledge, but no branch of knowledge (or combination of branches) can provide a foundation for logic. Yet, since a serious error in logic could undermine other branches of knowledge (and possibly our system of knowledge as a whole), a foundation for logic must be given. Must and cannot. Having postulated (i) that any resource for grounding logic must be more basic than the resources produced by logic itself, and (ii) that there are no resources more basic than those produced by logic, foundationalism rules out the existence of any rational resources for grounding logic.

This predicament leaves the logical foundationalist with two options: (i) show that logic does not require a foundation after all, or (ii) show that it is possible to provide a foundation for logic without using any resources produced by our system of knowledge. The first option, as we have seen in Chapter 2, Section 1, is a *reductio ad absurdum* of the foundationalist approach: if logic is ungrounded, then, due to its

position in the foundationalist hierarchy, our entire system of knowledge is ungrounded. The second option—*foundation without resources produced by our system of knowledge*—is the foundationalist's solution of choice. A foundation for logic, according to the foundationalist, must be free-standing—utilizing no knowledge-based resources whatsoever. Three contenders for a freestanding foundation of logic are *common-sense obviousness*, *conventionality*, and *pure intuition*. The first two, however, are highly problematic, and the third, too, at least in some of its forms, is precarious. A lengthy discussion of these solutions would distract us from our constructive goal. But briefly, let us note some of their difficulties.

II. Alleged remedies to the foundationalist predicament

(a) *Common-sense obviousness*. In conversation one often hears the view that logical theory requires no justification since logical truths are *obvious*, either in the sense that their content is trivial or in the sense that their truth is immediately clear to us.

The view that logic is justified by its (common-sense) obviousness is often traced to Quine, although at least in one place (1954a: 112) Quine says that this view is utterly unilluminating, likening it to the conventionalist view. Be that as it may, the idea that logical theory is justified by its obviousness is subject to multiple criticisms. They include:

(i) *Fallibility*. Judgments of obviousness are not just defeasible, they are highly fallible. As a result, a theory's correctness cannot be established solely, or even primarily, by appealing to its common-sense obviousness. Indeed, we may view the development of sophisticated, abstract, and far from common-sensical theories in all fields of knowledge as an antidote to the idea that common-sense obviousness guarantees veridicality. What appears at first obvious—that the Sun revolves around Earth, that the physical universe is Euclidean, that there is only one infinity, that whales are fish—turns out, upon critical examination, to be either false or far from obvious.

(ii) *Vagueness*. The idea that logic is a theory of obvious truths is a vague idea: Is logic a theory of *all* obvious truths, of *only* obvious truths? Of truths that are obvious *themselves* or are *obtained from* obvious truths? From obvious truths of *what kind*? (Any kind? Some *specifically logical* kind?) By steps of what kind? (Obvious steps? If so, obvious steps of *what type*?) By steps and truths obvious *to whom*? Obvious *in virtue of what*?

(iii) *Need for a Theoretical Explanation*. The vagueness of the idea that logic is obvious is not just inconvenient; it threatens the view that logic requires no theoretical foundation. To understand the sense in which logic is obvious and sanctioned by its obviousness, we have to understand (1) what features of logical truths/inferences make them obvious, (2) what features distinguish them from non-obvious truths/inferences and from obvious but not logical truths/inferences, (3) what the normative force of logical obviousness is, and so on. In short, to understand the view

that logic is grounded in the obvious we have to develop a *theory* of the foundational obviousness of logic. But this would belie the claim that (due to its obviousness) logic does not require a theoretical foundation.

(iv) *Low Standards of Knowledge*. As a standard of knowledge, common-sense obviousness is an exceedingly weak standard. Letting what appears to be the case (prior to critical examination) determine what logical theory to accept is utterly unreasonable. A foundationalist might respond by saying that there is no choice but to use low standards of knowledge for logic: either we ground logic in common-sense obviousness or we do not ground it at all. But this dilemma, as we have seen in Chapter 2, Section 2, is avoidable. We can replace the foundationalist methodology by a non-foundationalist yet foundational methodology: e.g., foundational holism.

(b) *Conventionality*. Conventionalism professes to ground logic in *linguistic convention*. A well-known advocate of this doctrine is Carnap (1939, 1950, and elsewhere).¹¹ Carnap's epistemic model divides our system of knowledge into two parts: *science* and *metascience*. Science is *reality-oriented*, metascience is *language-oriented*. Scientific knowledge requires a grounding in *experience*, metascientific knowledge requires a grounding in *pragmatic conventions*. Logic, as a discipline, belongs in metascience. Its role is not to obtain logical knowledge, but to *create conventional tools* for obtaining scientific knowledge. As such it vacuously meets the foundationalist requirement that "basic theories" be immune to error. To introduce error into our system of knowledge a theory has to *say* something, but logical theories say nothing; they merely install linguistic rules. Such rules, being conventional, are convenient or inconvenient, useful or useless, but they cannot conflict with, or be refuted by, reality.

Clever as this solution may sound, it, too, has serious (and for the most part well-known) drawbacks:

(i) *Trivialization of Knowledge*. To take a central branch of knowledge and declare it true by convention is to release it from the burden of veridicality, i.e., from the burdens of truth, evidence, and justification. This, in turn, is tantamount to stripping it of its epistemic entitlements. That point was made by Russell with respect to Poincaré's conventionalism and by Quine with respect to Carnap's. The method of conventional postulation, Russell says, is the "method of 'postulating' what we want". This method "has many advantages; they are the same as the advantages of theft over honest toil" (Russell 1919: 71).

¹¹ At least under certain interpretations inspired by Quine. These interpretations have been recently challenged by Friedman (1999) and others. One option is to attribute this view to "Quine's Carnap". For earlier discussions of Quine's Carnap see Chapter 3 and Chapter 6, Section 1.

Using a combination of a slippery-slope and *reductio-ad-absurdum* argument, Quine denounces the conventionalist characterization of logic as leading to a trivialization of our entire system of knowledge:

If... the [logical] primitives *can* be conventionally circumscribed in such fashion as to generate all and only the accepted truths of logic... the characterization is empty... [since] the same might be said of any body of doctrine as well (Quine 1935: 102).

For example:

This... method... is available... for any... branch of mathematics... In each case we merely set up a conjunction of postulates for that branch as true by fiat, as a conventional circumscription of the meanings of the constituent primitives, and all the theorems of the branch thereby become true by convention: the convention thus newly adopted together with the [logical] conventions... In this way mathematics becomes conventionally true... by proceeding from linguistic convention in the same way as does logic (Quine: 99–100).

But that is not all:

[This] method can... be carried... into the so-called empirical sciences. Having framed a maximum of definitions in the latter realm, we can circumscribe as many of our “empirical” primitives as we like by adding further conventions to the set adopted for logic and mathematics; a corresponding portion of “empirical” science then becomes conventionally true in precisely the [same] manner (Quine: 100).

The conventionalist construal of logic, Quine thus argues, makes a sham of knowledge. The conventionalist (in our terminology) exempts logic from the *friction* requirement, and in particular from *veridical friction*. In so doing, he is setting an exceedingly weak standard of knowledge for logic. If all a logical theory has to do to pass the epistemic bar is to have its author introduce it by postulation, its adherents declare it true by convention, or its supporters justify it as convenient, then the epistemic bar is set very low for it. And if the bar is set so low for one area of knowledge, why not for another? Why not for all?

(ii) *Qualitative Difference between Logical Laws and Conventions*. While few would disagree that conventionalizing the entire base of our system of knowledge would trivialize foundationalism, some might argue that conventionalizing certain circumscribed portions of the base, specifically *logic*, would be harmless. Including logical conventions in our system of knowledge, they might say, is as innocent as including any other conventions in it, for example abbreviations of the kind “1 kilometer = 1000 meters”. There is, however, an important difference between mere abbreviations and logical laws: the former are theoretically superfluous, while the latter are not. We can use logic to sanction abbreviations, but we cannot use abbreviations to sanction logic. Since logic’s correctness is crucial for every branch of knowledge, logic is saddled with a strong justificatory burden, one that is considerably greater than the burden of sanctioning abbreviations.

(iii) *Undermining Foundationalism from Within.* Considered as an internal solution to the foundationalist problem of logic, conventionalism violates one of the cornerstones of foundationalism: an absolute ban on circularity. The problem, as Quine (1935) explained, is that to regard the infinitely many (undecidable collection of) logical truths and consequences of a reasonably rich logical system as applications of a small number of general logical conventions (or at most a decidable set of logical conventions), logic itself must be used to justify the passage from these general conventions to their particular applications, e.g., through the use of Universal Instantiation. This stands in direct contradiction with foundationalism's demand that the resources used in justifying logic be more basic than, hence different from, those provided by logic itself.

(iv) *Introducing Errors into our System of Knowledge.* The conventionalist solution to the foundationalist problem of logic is often viewed as a "safe" solution. By treating logical laws as conventions we shield ourselves from *errors* in logic. If our logical "laws" make no factual claim whatsoever, they cannot be said to be false, or erroneous, or conflict with reality. The possibility of error in logic is an issue we will discuss in some detail in Chapter 10. Here it suffices to mention one widely known challenge to the "safety" of logical conventions—Prior's (1960) challenge. Suppose that we conventionally introduce a new logical constant, a binary connective, "tonk", into our logical system by two conventional "laws" that take the form of rules of inference: (1) $\Phi \vdash \Phi \text{ tonk } \Psi$, and (2) $\Phi \text{ tonk } \Psi \vdash \Psi$. Each "law" appears innocuous, but together they sanction the addition of any sentence whatsoever to our body of knowledge, including logical contradictions and sentences conflicting in a variety of ways with existent sentences.

The conventionalist might say that one can avoid tonk-like expressions by restricting oneself to *conservative* conventions, i.e., conventions that do not affect the "old" parts of our logical system (those that do not contain the new logical constants). This was Belnap's (1962) solution to the "tonk" problem. But the conventions the foundationalist is interested in have to ground our logical system in its *entirety*; so there is no preexisting logical subsystem that would constrain these conventions. And Belnap's solution, indeed, was not proposed as a defense of conventionalist foundationalism.

One could suggest that the "preexisting logical subsystem" is the informal stock of logical knowledge that has presumably guided us all along. This tactical move would not work either. If logic is grounded in informal logical knowledge—which, for the move to work, must be non-conventional—then logic is ultimately grounded in something other than convention.

(v) *The Unavoidability of Facts.* There is an important sense in which even the conventionalist is subject to factual standards or requirements. For example, no matter what conventions the conventionalist introduces as logical, he must subject these conventions to the requirement that inferences licensed by them *are in fact*

truth preserving, that these conventions in fact do not introduce fatal errors into an otherwise sound system of knowledge, that the resulting logic is in fact learnable and communicable, that it is in fact comprehensive enough to provide a logical framework for science and mathematics, etc. All these must be facts, not just convenient stories the conventionalist tells himself. But if conventions are subject to factual constraints, logic cannot be grounded only in convention.¹²

(c) *Pure intuition.* A Platonistic view, sometimes associated with Gödel, says that logic is grounded in a “second” reality, separate from the reality in which empirical science is grounded, and that humans access this reality through a special “additional sense”: “reason” or rational intuition (Gödel 1953–9: 354). Thus, “the proposition stating *modus ponens*”, Gödel suggests, is “an immediate datum” that “*can directly be perceived to be true*” (Gödel: 347 fn. and main text). To perceive *modus ponens* we use a sixth sense, a sense that grounds the laws of logic in much the same way that the five recognized senses ground the laws of science. One of the special features of this sixth sense is its ability to present to us highly “general (namely the ‘formal’) concepts and their relations” (Gödel: 354). These concepts and relations belong to a “second reality”, which is “completely separated from” space-time reality (Gödel: 353 fn.), though it is just as objective.¹³

Such a Platonist conception of logic, however, is problematic in much the same way that a radical Platonist conception of any abstract discipline is (see Chapter 5, Section 2). One of its problems is the application problem, i.e., the problem of why and how a logic grounded in the Platonic domain could work in the empirical domain. One may suggest that there are two logics—one for each domain, but this would raise the question: What is the logic of the empirical domain grounded in? And which logic is responsible for the logical truth and validity of mixed statements and inferences, those involving both abstract and experiential vocabulary, which are prevalent in science and everyday discourse?¹⁴ These problems are closely related to the bifurcation problem, which is in effect a *unity* problem. The foundationalist seeks to provide a unified foundation for human knowledge—one that encompasses both empirical and abstract knowledge—and an integral part of this goal is a unified

¹² We have discussed this point earlier in Chapter 8, Section 3 (see, in particular, the citation from Fine). Note that it would not do to say that the factual element is metatheoretical. It is essential for foundationalism that its foundation be comprehensive, i.e., that it encompass our system of knowledge in its entirety, including its metatheory.

¹³ It should be noted (1) that Gödel does not distinguish between logic and mathematics with respect to this view, and (2) that there are also pragmatist trends in Gödel’s philosophy of logic and mathematics. We may go further and suggest that our own conception of intellectual knowledge (as exemplified by the activity of “figuring out”, discussed in Chapter 5, Section 3) can be connected to some of Gödel’s ideas (including the combination of emphasis on intellectual discovery and recognition of pragmatic elements in theory development).

¹⁴ A similar problem was raised for truth pluralism by Tappolet. (See Chapter 7, Section 2 and Chapter 8, Section 4.)

logical method. But how can the foundationalist unify the logic of empirical knowledge with the logic of abstract knowledge?

These and other problems suggest that a rehabilitation of foundationalism based on a Platonist approach to logic is not promising.

I should emphasize once again, however, that not *all* elements of Platonism should be rejected. Many philosophers are drawn to Platonism for good reasons, like a *serious attitude* toward *abstract knowledge* and a *healthy critical attitude* toward *available alternatives* to Platonism, some of which, like conventionalism and extreme empiricism, are unreasonably limiting. The first point calls for further elaboration: Many philosophers recognize that abstract disciplines, like logic and mathematics, are in the business of discovering truths and constructing theories subject to high standards of knowledge, including high veridicality standards. This recognition is supported by the existence of highly developed disciplines (such as logic and mathematics) whose accomplishments arguably include the discovery, or account, of laws governing such features. This said, the Platonist solution to the foundational problem of logic is not the only possible solution having these advantages; nor is it the best solution.¹⁵

A common feature of the above “remedies” to the basic-knowledge predicament (within the foundationalist ideology) is that they all lack adequate friction. Platonism, conventionalism, and common-sense obviousness all exemplify Kant’s dove who, by looking for an easy way to reach her destination, eliminates the very condition that makes it possible for her to reach it. Indeed, in the case of a foundational project, friction is not just a means but also an integral part of the destination.

9.3 Solution: The Foundational-Holistic Methodology

The main obstacle to constructing a philosophical foundation for logic is the deeply ingrained (though often tacit) tendency to view such a foundation through the *foundationalist* lens. A natural way to remove this obstacle is, therefore, to replace this lens. The *foundational-holistic* methodology developed in Chapter 2, Section 2 of this essay offers a different template for a foundational project. This methodology is expressly designed to encompass all branches of knowledge, including highly abstract and “basic” branches, such as logic, making a foundation for such disciplines possible.

One characteristic of the foundational-holistic methodology is its serious attitude toward highly abstract disciplines. Like Platonism, it regards such disciplines as genuine fields of knowledge, and it affirms the universal grounding requirement as it applies to them. At the same time, foundational-holism rejects the unreasonable demands attached to the grounding requirement by the foundationalist methodology, in particular, strict ordering and absolute certainty. In providing a grounding for abstract fields of knowledge, foundational holism says, we should not burden

¹⁵ I should add that contemporary construals of intuition as a tool of discovery get further and further away from Platonist conceptions of intuition. See, for example, Parsons (1995, 2000, 2008).

ourselves with superfluous restrictions. It insists on robust veridicality standards for such disciplines, but it does not impose either a specific structure or limited resources on the grounding process. Foundational holism's motto is "friction through freedom": freedom in seeking appropriate ways of satisfying the friction requirement. In accordance with the "Neurath boat" metaphor, foundational holism views the grounding process as an ongoing, step-by-step process, one that uses whatever resources are available *here and now* to ground each discipline. Starting with a partial foundation for logic, one creates additional resources for strengthening, extending, and/or revising this foundation, arriving at a more complete and robust foundation as one goes.

Moreover, foundational holism frees us to use resources created by other disciplines to ground logic without worrying about (non-destructive) circularity, as explained in Chapter 2, Section 2. The grounding project is not an "either-or" project: either we provide an ultimate, eternal foundation for logic, or we make do with a purely pragmatic justification of our logical theory (or exempt it from grounding altogether). The unfeasibility of an absolute foundation for logic does not mean that logic gets a *free pass*. As in the case of empirical science, there is a serious question of rightness and error in logic—not just error in *applying* the logical laws, but error in determining what these laws *are*.

Using a variation on Sellars (1956: 78) we may express our view as follows:

To accept that logical knowledge has *no* foundation is to put it in a box with *rumors and hoaxes*.¹⁶

And in the same spirit (see Sellars: 78–9) we may add:

Above all, the foundationalist picture is misleading because of its static character. For logical knowledge is rational, not because it has an *absolute* foundation but because it is a *self-correcting enterprise* which can and must put *any* claim in jeopardy, though not *all* at once.¹⁷

Our task, then, is to provide a foundation for logic in the sense of a *critique* of logic: to challenge our logical theory with critical questions (in particular veridicality questions), to develop a method for constructively answering these questions, and to use this method to answer these questions. The special difficulty in grounding logic, as distinct from other disciplines, lies in the combination of abstractness and basicness. By the basicness of logic I mean *not* its basicness in the foundationalist sense, *but* the fact that logic is incorporated in all other branches of knowledge, that we cannot make any significant step in either practical or theoretical cognition

¹⁶ Sellars: "to say that empirical knowledge has *no* foundation . . . is . . . to put it in a box with rumors and hoaxes".

¹⁷ Sellars:

Above all, the picture is misleading because of its static character. . . . For empirical knowledge, like its sophisticated extension, science, is rational, not because it has a *foundation* but because it is a self-correcting enterprise which can put *any* claim in jeopardy, though not *all* at once.

without using logic. Foundational holism enables us to deal with this complication by recognizing the existence, and sanctioning the joint use, of two multi-patterned and interconnected networks: a network of connections *among units of knowledge*, and a network of connections *between units of knowledge and reality*. The existence, in principle, of multiple cognitive routes to reality opens up the possibility of reaching diverse aspects of reality, including those aspects which are central to a grounding of logic. And the existence of interconnections among multiple units of knowledge means that in grounding logic we are free to avail ourselves—though, as always, critically—of resources created by disciplines which are less basic than, or as basic as, logic (in the above sense). And while foundational holism does not require either an absolute foundation for logic or a strictly linear foundation, it requires a *solid* theoretical foundation and provides tools for constructing one.

With this workable methodology at hand, we are ready to proceed to the second part of our solution to the foundational problem of logic: the construction of an actual foundation for logic, or an outline of such a foundation. In attending to this task I will treat the foundational problem of logic as a *theoretical problem* rather than as a problem of spelling out our *pretheoretical intuitions*, and as an epistemic problem rather than as a problem of making sense of natural-language discourse.

10

An Outline of a Foundation for Logic

Although it is clear that the foundational-holistic methodology makes a foundation for logic possible, it is far from clear how to go about constructing such a foundation. What investigations shall we start with? What question is the key to discovering what we are seeking? For me, this conundrum took the form of a metaphor—entering the old city of Jerusalem. There are many gates leading into the city and it does not matter what gate we enter through. In the same way, I thought, there are many questions that can lead to a foundation for logic; all one needs is one question that will work. So I looked for one question that would be (i) new and truly open for me (so I would approach it without preconceptions), (ii) potentially connected to the nature of logic, (iii) sufficiently precise to allow a definite investigation, and (iv) philosophically rich and stimulating. The question I found was: what is the nature and scope of logical constants?, or more precisely: is there a philosophically well-motivated criterion for logical constants? Now, some philosophers regard the question of logical constants as tangential to the foundational problem of logic. And it is possible, indeed, to approach this question from directions that do not lead to foundational results. But, given the central role of logical constants in determining logical truths, logical consequences, logical contradictions, etc., it is also possible to approach it in a way that does. My own way of approaching this question was *functional*. I asked: *if the function or task of logic is X, what choice of logical constants will enable it to accomplish X, why, and how?* (and I used Tarski's 1936a intuitive characterization of logical consequence as a stepping stone for determining what X is). Following my investigations of this question, my inquiry continued in a holistic, back-and-forth manner, adding new questions and proceeding along what I now call foundational holistic lines.

In presenting my outline of a foundation for logic in this essay, however, I will follow the order of explanation and justification rather than the order of discovery. But the “functional” structure of the investigation will be preserved, and indeed extended. The underlying questions will have the form: “If the function (goal, task, role) of logic in our system of knowledge is X, what kind of grounding will enable it to accomplish X?”.

10.1 What is Logic's Task in our System of Knowledge?

What kind of discipline is logic? Is it a genuine branch of theoretical knowledge or just an instrument of knowledge? Many believe that logic is not a genuine branch of knowledge, and one common argument for this view appeals to its *topic neutrality*. For a field to be topic neutral, so the argument goes, is for it to have no topic of its own, that is, no topic or subject matter for a theoretical study. Physics, mathematics, psychology have distinct subject matters, of which they provide (or at least seek to provide) theoretical knowledge. But logic does not.

This view is based on a confusion. Logic is indeed topic neutral, but being topic neutral is not the same thing as not having a subject matter of its own. Logic does have a subject matter of its own. Its subject matter is logical inference, logical truth, logical contradiction (inconsistency), logical equivalence, etc., where these are very different subject matters from those of physics, mathematics, or psychology. In spite of having a definite subject matter, however, logic is topic neutral. Its topic neutrality consists in the fact that it *applies the same tests* of logical validity, logical truth, etc., to inferences and sentences in *all* area of discourse, regardless of *their* subject matter. Logic, thus, is a theoretical discipline with its own subject matter, and one of its jobs is to provide theoretical knowledge about it. The theoretical yield of logic is further extended by metalogic. In the last century, metalogic produced numerous theoretical results: completeness (of first-order logic), incompleteness (of first-order arithmetic, many other parts of mathematics, second-order logic), the Löwenheim-Skolem-Tarski results, Church's thesis, the undecidability of the halting problem, and so on, which throw important light on logic, and many of which have significant consequences for our system of knowledge as a whole. Moreover, due to a host of connections between the subject matter of logic and those of other disciplines, logic can, and has, joined forces with linguistics, computer science, mathematics, philosophy, etc., to provide new knowledge on a large array of interdisciplinary topics as well as on specific topics within some of these disciplines.

But being a branch of knowledge does not mean that logic is not also an *instrument* of knowledge. An examination of logic's historical role in our system of knowledge shows that it plays both roles. On the one hand, logic provides theoretical knowledge of its subject matter; on the other hand, it provides a method for expanding knowledge and detecting (preventing, correcting) error in all fields of knowledge. Its two roles are closely connected: the method generated by logic is based on its study of logical inference, logical contradiction, etc., and one of the things that logic (as a branch of knowledge), and in particular metalogic, studies is the logical method. In using logic's role in knowledge as a key to figuring out its foundation, I will focus on its instrumental role.

To understand this role, let us begin with the straightforward observation (appealed to earlier on) that our biological, psychological, and other cognitive limitations considerably restrict our ability to acquire knowledge. A significant aspect

of this handicap is the severe limit it sets on our ability to attain *direct* knowledge. This severe limitation introduces an urgent need for *indirect* ways of expanding our knowledge and detecting error in our thought (theories). Now, given that knowledge qua knowledge is knowledge of the world and that as such its first and foremost (though not only) standard of success is truth, one effective means of expanding knowledge is the method of *inference*, i.e., the method of transmitting truth from sentences to sentences based on some of their features. And a method of inference that is *highly general* and *modally powerful* is especially valuable. The same holds for methods for identifying errors. A method that is not restricted to a specific area of knowledge and identifies an especially pernicious type of error, i.e., an error that has the potential of undermining an entire theory and even our entire body of knowledge, is particularly desirable.

The task of logic in our system of knowledge is the development of just such a method. Focusing on inference, we can express this by the thesis:

Thesis 1: Function (Task) of Logic

Logic's task is to develop a method of inference which is both highly general and has an especially strong modal force. More specifically, its task is to develop a method for constructing inferences that transmit truth from sentences to sentences with an especially strong modal force and regardless of field of knowledge.

Our strategy is to examine logic in light of this task. We ask two questions: (i) What kind of foundation (grounding) is *required* for logic to perform this task? (ii) What kind of foundation will *enable* it to perform this task? In asking these questions, we keep away from the traditional quest for a “necessary and sufficient condition”. Approaching the foundation of logic holistically, we think of these questions in terms of “constraints” (things that need to be, or not to be, the case in order for logic to reach its goal) and “opportunities” (things that would enable logic to reach its goal). And we think of an adequate foundation for logic as addressing both questions.

One basic aspect of our question concerns what kind of thing logic is grounded in: mind, world, something else? Our first question is, therefore, whether logic is grounded in the human mind (language, concepts, conventions, etc.) or in the world (certain aspects of reality).¹

10.2 Is Logic Grounded in the Mind or in the World?

The prevalent view today is, as it has been for centuries, that logic is grounded *solely* in the mind: in the structure of our brain, in the rules of our language, in our psychological makeup, in our pragmatic decisions, in the form of human thought,

¹ Due to the functional nature of our investigation, the questions “What *is* logic grounded in?” and “What *should* (*ought*) logic be grounded in?” coincide. Accordingly, I freely switch between these two forms of speech.

or in something else of a similar kind. Most philosophers believe that logic has everything to do with the mind and nothing to do with the world. A traditional representative of this view is Kant:²

[P]ure general logic... deals with **nothing** but the mere **form of thought** (Kant 1781/7: A54/B78).³

This renders logic *analytic* in Kant's terminology.

A more recent view grounds logic in the mind in the psychological sense. Hanna (2006), for example, develops an account of logic that grounds it exclusively in human psychology: logic is created by rational animals based on an innate template, called *protologic*, which belongs to a special cognitive faculty, the *logic faculty*, and the study of logic and its special faculty is a common project of cognitive psychology and philosophy. Maddy (2007) offers a strongly naturalistic account of logic's grounding in the mind, combining elements from Kant and contemporary psychology (e.g., Spelke et al. 1995). Her correlate of Hanna's protologic is "rudimentary logic".⁴

Alongside the majority view there is also a minority view which says that logic is grounded in the world, albeit one that has rarely been developed systematically or in much detail. An early representative of this view is Wittgenstein in the *Tractatus* (1921). Although Wittgenstein does not believe that a *theoretical* grounding of logic is possible, he thinks that logic *is* grounded in reality in the sense that reality has a logical structure and that the logical structure of our language reflects this structure. Thus he says:

Propositions *show* the logical form of **reality**. They display it. (Wittgenstein 1921: 4.121).

The fact that the propositions of logic are tautologies *shows* the formal—logical—properties of language and the **world** (Wittgenstein: 6.12).

The propositions of logic describe the scaffolding of the **world**, or rather they represent it (Wittgenstein: 6.124).

Logic is... a mirror-image of the **world** (Wittgenstein: 6.13).

A second proponent of realism with respect to logic is Russell (in some of his writings). In his 1919 book Russell famously says that "logic is concerned with the **real world** just as truly as zoology, though with its more abstract and general features" (Russell 1919: 169).

A later supporter, or rather precursor, of the view that logic is grounded in reality is Tarski. Although Tarski himself never explicitly addressed the relation between logic and reality in his theoretical writings, some of his views are strongly suggestive of a significant logic–world connection. Consider, in particular, Tarski's two classical

² For **boldface** within citations see fn. 3, Chapter 1.

³ For additional citations see Kant (1781/7: Axiv, Bix, A52–5/B76–9, A130–1/B169–70, A796/B824, A796/B824), and Kant (1770s–1800).

⁴ It should be noted, however, that Maddy's grounding of logic is also naturally viewed as involving the world (see p. 258).

papers, “The Concept of Truth in Formalized Languages” (1933) and “On the Concept of Logical Consequence” (1936a). In his 1933 paper Tarski characterizes *semantic concepts* as concepts that have to do with the *relation between language and reality*, hence as concepts that significantly involve *reality*:

A characteristic feature of the **semantical** concepts is that they give expression to certain relations between the expressions of language and the **objects about which these expressions speak** (Tarski 1933: 252).

Essentially the same characterization appears in Tarski (1936b):

We shall understand by **semantics** the totality of considerations concerning those concepts which, roughly speaking, express certain connexions between the expressions of a language and the **objects and states of affairs** referred to by these expressions. As typical examples of semantical concepts we may mention the concepts of *denotation*, *satisfaction*, and *definition* [*i.e., unique determination of objects*]. . . . The concept of *truth* also . . . is to be included here, at least in its classical interpretation, according to which ‘true’ signifies the same as ‘corresponding with **reality**’ (Tarski 1936b: 401).

Now, in his 1936a paper Tarski sets out to provide a *semantic* definition of logical consequence, stressing that it is the *semantic* character of this definition that is the key to its success. Specifically, Tarski defines logical consequence in terms of “model of”, where “model of” is defined in terms of *satisfaction*, a classical *semantic* notion in Tarski’s sense, *i.e.*, a notion that has to do with the relation between language and objects (states of affairs). Putting the two together, we see that Tarski chose to define logical consequence in terms that significantly involve *world* (objects). Whether Tarski himself recognized the connection between what he said about semantic notions in general and his choice of a definition of logical consequence I do not know, but this connection is there for all to see. The fruitfulness of Tarski’s semantic approach to logical consequence suggests that we should not take it for granted that logic is grounded exclusively in concepts (language, mind), but seriously look into the possibility that it is grounded in objects (world) as well. Tarski’s letter to Morton White concerning revision in logic explicitly suggests that “certain . . . very fundamental” features of the world might be involved in such revision (Tarski 1944a: 31–2).⁵

A fourth philosopher who connects logic to reality is Quine (in some of his moods). As we recall from Chapter 6, Section 1, Quine declares that “[l]ogical theory, despite its heavy dependence on talk of language, is . . . **world-oriented** rather than language-oriented”. He then goes on to say that it is “the **truth** predicate” that “makes it so” (Quine 1970/86: 97). Although Quine himself is ambivalent about this observation, his cryptic explanation, as we shall soon see, is right to the point.

Another supporter of the view that logic is grounded in reality is the early Putnam. According to Putnam, just as relativity theory teaches us that we live in a world with a

⁵ I will return to this letter in Section 10.9.

non-classical (non-Euclidean) geometry, so quantum mechanics teaches us that “[w]e live in a **world** with a non-classical logic” (Putnam 1968: 184). More recently, Almog has said that “the necessity of logical truths won’t come from a trivial observation on the meaning of ‘logical truth’” and that “logic focuses on special, structural (‘formal’), traits of **actuality**” (Almog 1989: 214). Brandom also seems to support some version of this view. Speaking about inference, he says that to be *objective*, an inference—including a logical inference—requires something like a grounding in reality:

It must be shown how . . . it is possible to fund *objective* properties of inferring[.] . . . [W]hat it is correct to conclude or to say depends on how the **objects** referred to, talked about, or represented **actually are**. . . . [N]ot just any notion of correctness of inference will do . . . A semantically adequate notion of correct inference must . . . fund the idea of *objective* truth conditions and so of *objectively* correct inferences. Such proprieties of . . . inference . . . are determined by **how things actually are**, independently of how they are taken to be. Our cognitive attitudes must ultimately answer to these attitude-transcendent **facts** (Brandom 1994a: 136–7; see also pp. 78, 280, 649).

I have argued for this view in many of my publications, starting with Sher (1989). Shapiro, arguing against Resnik’s logical irrealism, also emphasizes the *objectivity* of logic and the need to explain it in a way that fits in with a realist conception of reality:

[T]he central role of logic is not explicable solely in terms of ‘our linguistic and inferential convention’. A common slogan is that logical consequence is *truth preserving*. The relevant feature here is that consequence must preserve *objective truth*, since logic is applicable in areas which are up for **realist** construal, such as science. Logical consequence must answer to, and preserve, **objective truth**. That’s the point of logic. So **logic is not completely dependent on just our conventions** . . .

Suppose, for example, that a community adopted a tonk-like connective. Then their consequence relation would be faulty, on **objective** grounds, simply because it fails to preserve truth. . . .

[A] lot of **objectivity** must be built into logical consequence . . . Logic is **objective** if anything is (Shapiro 2000: 354–5).

Chateaubriand, too, is very clear about logic being “a theory of **reality**” (Chateaubriand 2001: 17): “[L]ogic is based on considerations about truth and **reality**” (Chateaubriand: 21), he says; “the laws of logic . . . express fundamental features of the structure of **reality**” (Chateaubriand: 29).⁶ Similarly: Priest’s answer to the question “should we be a realist about logic?” is “yes” (Priest 2001: 36).

Maddy, who advocates a strongly naturalistic approach to logic, says:

As a first approximation . . . the Second [i.e., strongly naturalist] Philosopher hopes to develop an account of logical truth [which includes the principle that] . . . logic is true of the **world** because of its underlying structural features (Maddy 2007: 226).

⁶ See also Chateaubriand (2013).

Goldfarb attributes the view that logic is world-oriented to Frege:

On Frege's *universalist conception* . . . the concern of logic is the articulation and proof of logical laws, which are universal truths.

For Frege, the laws of logic . . . state truths about everything.

[O]n the universalist conception logic . . . issu[es] laws that are . . . statements about the world (Goldfarb 2010: 68–9).

And Sider says: “Against logical conventionalism, I uphold Russell’s (1919, p. 169) diametrically opposed position” [cited above] (Sider 2011: 98), and: “Logical content is as ‘worldly’ as nonlogical content” (Sider: 276).

In spite of the fact that these and possibly some other philosophers claim, or suggest, that logic is grounded in the world, this view is still a minority view. Our own thesis is:

Thesis 2: Mind and World

Logic is grounded (requires a grounding) both in the mind and in the world, and these grounds are interconnected.

Another way to express this thesis is to say that logic lies both in the center and in the periphery, and these two locations are interrelated. This thesis is motivated both by general considerations, considerations pertaining to all branches of knowledge, and by specific considerations—ones specific to logic. The general considerations we have discussed in Parts I and II of this essay. The main point is that knowledge qua knowledge must be grounded both in its object, the world in a broad sense, and in features of its agent, the mind. Groundedness in the world is veridicality, i.e., compliance with strict standards of truth, evidence, and factual justification. Groundedness in the mind is groundedness in human cognitive capacities, both active and passive.⁷ The requirement that we subject logic not just to pragmatic standards but also to veridical standards is also an active constraint issued by the mind.

There is much to say about logic’s grounding in the mind, and in future work I hope to discuss this topic in greater detail. But the view that logic is grounded in the mind is already so deeply ingrained that it does not require new support. What does require new support, as well as explanation, justification, and systematic presentation, is the view that logic is grounded in reality.

To prevent misunderstanding, however, let me clarify at the outset that by saying that logic is grounded in the world I do not mean that logic is contingent in the sense that empirical science is. Nor do I mean that logic is grounded in anything like the “thing in itself” or “transcendent reality” in Kant’s sense. Likewise, my view is neither Platonist nor nominalist. Neither the Kantian duality of noumena and phenomena,

⁷ As noted earlier, a remarkable example of an active contribution to logic by the mind is Frege’s construction (creation, invention) of a new, symbolic language, *Begriffsschrift*, for the formulation and examination of logical proofs.

nor pure Platonism or extreme empiricism are compatible with my outlook. My claim that logic is grounded in the world is not deflationist either. Logic is grounded in the world in a substantive sense, a sense that yields substantive explanations, solves significant problems, and has non-trivial consequences, as we will in the following sections. Finally, let me emphasize that in saying that logic is grounded in the world I do not mean that it is grounded in all (or just any) features of the world. My claim is that there are certain highly specific features or aspects of reality that logic is grounded in, that its grounding in these features explains why it has the traits that it has (e.g., a strong modal force), and that the foundational philosopher's task is to *uncover* these features, explain how being grounded in these features enables logic to perform its task, and point out some significant consequences of this grounding.

10.3 Why does Logic Require a Grounding in the World?

What are the special reasons—reasons having to do with logic itself rather than with knowledge in general—that require logic to be grounded in reality (lie in the periphery)? Below I will focus on three reasons: a common-sensical reason, an example-based, partly historical reason, and a theoretical, systematic reason.

I. Logic has to “work” in the world

There is a very real, straightforward sense in which logic has to work in the world. Consider the task of building an airplane. To the same extent that to build an airplane we need an aerodynamic theory that actually works in the world, so to build an airplane we need a logical theory that actually works in the world. Put negatively: just as defective aerodynamical principles can cause an airplane to malfunction, so defective logical principles can result in its malfunctioning. If, in designing an airplane, we rely on incorrect logical laws—e.g., the law of “affirming the consequent”, or the “new Leibniz law” (see Subsection II)—we are likely to cause *drag* when *lift* is needed, a right turn when a left is intended, etc.⁸ A flawed logic can wreak havoc on an airplane no less than a flawed aerodynamics. And in general, employing a faulty logic—a logic with laws such as *affirming the consequent*—is likely to cause disruption in real-world enterprises—cause atomic plants to explode, cars to collide, workers to lose their salaries, and so on. Why is *affirming the consequent* likely to cause these things while *affirming the antecedent* (modus ponens) is not? Clearly, the reason is *not* that *affirming the antecedent* is more elegant, or more intuitive, or more natural to humans than *affirming the consequent*. The reason is that (under normal conditions) affirming the antecedent is in some important sense *in sync* with the world, while affirming the consequent is not. Moreover, since logic,

⁸ See Sher (2002).

unlike aerodynamics, is used in *all* (or most) areas of our life, a faulty logic would cause more damage than a faulty aerodynamics, and in this sense it is *more important* that our logical theory works in the world than that our aerodynamical theory does. But to work in the world, a logical theory has to be grounded in the world, at least in the negative sense of being *constrained* by the world.

Which features of the world ground logic is a question we will turn to in the next section. But it is clear that some rules which syntactically look like logical rules work in the world, while others do not. For that reason, we cannot take it for granted that any logical system we postulate, or any logical system that seems natural to us, works in the world. Indeed, what scientific realists say about scientific theories (and in particular about the abstract parts of such theories) applies to logical theories as well: it would be a mystery that a logical theory worked in the world (in flying airplanes, running atomic plants, etc.) if it were not in tune with the world. In designing a logical system, therefore, the world must be taken into account. This does not mean that there is no room to maneuver in designing such a system, or that all features of the world must be taken into account. But working in the world is a significant constraint on the choice, and construction, of a logical theory.

II. Lessons from history and hypothetical situations

The claim that logic is constrained by the world is based not just on common-sense considerations concerning logic in general, but also on historical and artificial examples of specific factual errors in logic.⁹ A dramatic historical example of this kind is the error in Frege's logic discovered by Russell. Frege's logic, Russell showed, is committed to the *existence* of an *object*, a class, that *does not, and cannot, exist*.¹⁰ This generated a fatal paradox that forced logicians to reject Frege's logic, at least in its original form.

Less conclusive, but still instructive, are cases where philosophers (or others) have claimed that "standard" logic (usually understood as the bivalent system of first-order mathematical logic) fails to work under certain conditions. Such claims have been made by proponents of free logics, fuzzy logic, quantum logic, and many other non-classical logics. Although some of these logics are controversial, they have numerous adherents and have received serious attention from the philosophical community as a whole. This suggests that there is a genuine question whether a given logical system is factually correct for, or leads to errors in, different domains.

Finally, it is easy to construct artificial examples of factually erroneous logical laws. Consider, for example, the introduction of the law " $\Phi(x); x \neq y \vdash \sim \Phi(y)$ " into a logical theory. The syntactic form of this law is very similar to that of Leibniz's law, and in this respect it appears unobjectionable. But for this law to be valid, a very special

⁹ Note: "factual", here, is synonymous with "having to do with reality", and like the latter, does not imply "contingent" or "empirical".

¹⁰ This is not the only possible construal of Frege's error, but it is a fairly straightforward construal.

condition has to hold in the world: objects in the world must have no common properties (or at least no common definable properties). This, however, is clearly not the case. The “new” Leibniz law, thus, is *factually* invalid. There might, of course, be other reasons for rejecting this law, but its conflict with reality is by itself a sufficient reason for rejecting it. Indeed, the factual conflict in this case is so deep that it renders this law classically inconsistent.

These and other examples of factual errors in logic—both historical and artificial—offer an additional reason for viewing agreement with reality as a major factor in the acceptance or rejection of a proposed logical system. The next consideration differs from the first two in being theoretical and systematic.

III. Logic is both constrained and enabled by reality due to and through its inherent connection with truth

An important consideration connecting logic to reality has to do with its inherent connection to *truth*. To explain this consideration we will focus on *logical consequence* (a semantic rendering of logical inference). Our discussion will be divided into two parts: A. Why and how logic is *constrained* by the world. B. Why, and how, logic is (can be) *enabled* by the world. Speaking in terms of “periphery”, we will explain why, given its special connection with truth, logic must lie in the periphery and how lying in the periphery enables it to perform its task.

A. *Why and How Logic is Constrained by the World.* Logic’s task, according to Thesis 1, is to develop an especially powerful method of inference, i.e., a method that generates (identifies) inferences or consequences which guarantee the transmission of truth from premises to conclusion with an especially strong modal force and based on universal principles (i.e., principles that apply in all fields of knowledge). This task, as we shall presently see, sets significant constraints on logic, in particular, constraints having to do with the world. To see what these constraints are, let us begin with consequences in general.

A consequence in general is a binary relation between a set of sentences and a sentence satisfying the following condition:

- (C) Sentence σ is a *consequence* of set of sentences Σ iff the *truth* of the sentences of Σ is transmitted to, or preserved by, σ (i.e., if all the sentences of Σ are true, so is σ).

One important distinction between different types of consequence relations satisfying (C) concerns their *modal force*, i.e., the force with which the truth of the premises (sentences of Σ) is transmitted to the conclusion (σ).¹¹ For the purpose of the present

¹¹ See the earlier discussion of modal force in Chapter 5, Section 3 and Chapter 8, Section 4. One way to express modal force is in terms of “possible worlds”, including “actual world”. Although in the main text I will not use the terminology of “possible worlds”, I will do so in footnotes.

investigation it will suffice to consider three types of consequence: *material* consequence, *nomic* consequence, and *logical* consequence.¹² We will characterize these as follows:

- (MC) σ is a *material* consequence of Σ iff the *truth* of the sentences of Σ is *materially* transmitted to, or preserved by, σ .¹³

An example of a material consequence is:

- (1) Barack Obama is president; therefore, Earth has one moon.

Turning to *nomic* consequence, we have:

- (NC) σ is a *nomic* consequence of Σ iff the *truth* of the sentences of Σ is *nomically* transmitted to, or preserved by, σ .

That is, in the case of nomic consequence truth is transmitted from Σ to σ with nomic force, i.e., the force of a law of nature.¹⁴ An example of a putative nomic consequence (-schema) is:

- (2) The force exerted by a on b is c ; therefore, the force exerted by b on a is c .

Turning to *logical* consequence, we are still not in a position to identify the exact type of modal force characteristic of this type of consequence. But we can start by saying:

- (LC) σ is a *logical* consequence of Σ iff the *truth* of the sentences of Σ is transmitted to, or preserved by, σ with an *especially strong* modal force, one that is stronger than the modal force of a law of nature.¹⁵

An initial paradigm of a logical-consequence schema is “Something is P and Q ; therefore, Something is Q ”. In symbols:

- (3) $(\exists x)(Px \ \& \ Qx)$; therefore, $(\exists x)Qx$.

¹² In Sher (1996) I discuss metaphysical consequence as well, suggesting that the modal force of most metaphysical consequences lies in between those of nomic and logical consequences.

¹³ Using the “actual world” terminology, we can restate (MC) in the form:

(MC*) σ is a *material* consequence of Σ iff in the *actual world* it is not the case that all the sentences of Σ are true and σ is false.

¹⁴ Using “possible-worlds” terminology, we have:

(NC*) σ is a *nomic* consequence of Σ iff there is no *physically possible world* in which all the sentences of Σ are true and σ is false.

¹⁵ Assuming that “!” indicates an especially broad type of possibility (the correlate of an especially strong modal force), we have:

(LC*) σ is a *logical* consequence of Σ iff there is no *!-possible world* in which all the sentences of Σ are true and σ is false.

Now, our first claim is that consequence relations in general—i.e., relations satisfying (C)—are constrained by the world. One line of reasoning leading to this conclusion is the following:

1. It is *inherent* in consequence relations that they hold between sentences due to certain connections between their *truth values*. (Consequence relations are *inherently* relations of transmission or preservation of *truth*.)
2. The *truth value* of a sentence *inherently* depends on how things are in the *world*. (A sentence is *true* iff *things* are as the sentence says.)¹⁶
3. *Therefore*, consequence relations, including the relation of *logical consequence*, have to respect, or take into account, the dependence of truth on the world. Specifically: *consequence relations must take into account the relations between the conditions that have to hold in the world for the sentences of Σ to be true and the conditions that have to hold in the world for σ to be true.*¹⁷

To see how we arrive at this line of reasoning as it pertains to *logic*, suppose someone comes up with a proposal for a logical theory, or system, \mathcal{L} . Now ask: Under what conditions is \mathcal{L} an acceptable logical theory, one that makes correct judgments about logical consequence? In particular, under what condition does \mathcal{L} makes correct judgments about whether the truth of given premises guarantees the truth of a purported conclusion with a modal force commensurate with logic's "job description"?

For the sake of simplicity let us consider a case in which \mathcal{L} claims that a sentence S_2 is a logical consequence of a sentence S_1 , where these sentences have distinct, non-controversial correspondence truth-conditions, i.e., truth conditions referring to the world.¹⁸ In symbols:

$$(4) \quad (\text{Level of Consequence}) \quad S_1 \models_L S_2.^{19}$$

Assume S_1 is true. Then, for (4) to be true, the truth of S_1 must guarantee the truth of S_2 with a certain modal force. I.e., it must be the case that:

$$(5) \quad (\text{Level of Truth}) \quad T(S_1) \text{ -----} \rightarrow T(S_2),^{20}$$

¹⁶ See Chapter 8, Sections 1 and 3.

¹⁷ It is significant that this argument is based on *inherent*, rather than *accidental*, features of consequence and truth.

¹⁸ I will discuss the use of logic in contexts not involving truth (or correspondence truth) in Section 10.10.

¹⁹ Notation: " $X \models Y$ " means: " Y is a consequence of X (of an unspecified kind)". " $X \models_L Y$ " means " Y is a *logical* consequence of X ".

²⁰ Here and below different styles of arrow are used to represent different connections. Where the types of connection are not clear from the context, they will be indicated either in the text or in footnotes.

where the arrow indicates a guarantee of transmission of truth with the requisite modal force. Now, let \mathfrak{C}_1 and \mathfrak{C}_2 be the conditions (situations) required to hold in the world so that S_1 and S_2 are true:²¹

$$\begin{array}{ccc} (6) & (\text{Level of Truth}) & \begin{array}{cc} T(S_1) & T(S_2) \\ \Downarrow & \Downarrow \\ (\text{Level of World}) & \mathfrak{C}_1 & \mathfrak{C}_2. \end{array} \end{array}$$

Then, there are at least four situations in which *the world can falsify* (4)—four ways in which the world can falsify \mathcal{L} 's claim that S_2 is a logical consequence of S_1 , or four ways in which *the world constrains our logical theory*:²²

a. Assume the world is such that \mathfrak{C}_1 being the case rules out \mathfrak{C}_2 being the case, i.e., the world being the way it has to be in order for S_1 to be true *rules out* its being the way it has to be in order for S_2 to be true:

$$(7) \quad (\text{Level of World}) \quad \llbracket \mathfrak{C}_1 \rightarrow \text{NOT } \mathfrak{C}_2 \rrbracket.$$

This is sufficient to *falsify* (4), i.e., to falsify \mathcal{L} 's claim that S_2 is a logical consequence—or indeed, a consequence of any kind (material, nomic, or *logical*)—of S_1 . I.e.:

$$(8) \quad (\text{Level of Consequence}) \quad \text{NOT: } S_1 \models_{M,N,L} S_2.^{23}$$

No matter what the proposed logical theory, \mathcal{L} , says, S_2 is *not* a logical consequence of S_1 .

b. Now suppose that the world does not rule out that things are the way both S_1 and S_2 require them to be in order to be true, but that as a *matter of accident* things in the world are as S_1 requires but not as S_2 requires:

$$(9) \quad (\text{Level of World}) \quad \llbracket \mathfrak{C}_1, \text{NOT } \mathfrak{C}_2 \rrbracket.$$

This, too, is sufficient to *falsify* \mathcal{L} 's claim that S_2 is a logical consequence—or a consequence of any kind (material, nomic, or *logical*)—of S_1 :

$$(10) \quad (\text{Level of Consequence}) \quad \text{NOT: } S_1 \models_{M,N,L} S_2.$$

²¹ Note: throughout this chapter we will alternate between describing \mathfrak{C} as a *condition* and a *situation*: \mathfrak{C} is the *condition* that has to be satisfied by the world for a given sentence (set of sentences) to be true (i.e., \mathfrak{C} is the *truth condition* of that sentence) or, alternatively, \mathfrak{C} is the *situation* that has to hold in the world for this sentence to be true and is responsible (on the world's part) for its truth, if it is true.

²² (i) In speaking about “situations”, “facts”, and “states of affairs” in this essay I am speaking informally, in a way that does not involve any commitment to a particular kind of objects or ontology.

(ii) The four ways in question are not mutually exclusive. I allow redundancy in order to introduce pertinent distinctions.

²³ The subscripts “M” and “N” stand for “material” and “nomic” respectively.

c. Next suppose that the world is such as to make both S_1 and S_2 true, but the world being as S_1 requires *does not guarantee* (with a non-negligible modal force) its being as S_2 requires:

$$(11) \quad (\text{Level of World}) \quad \llbracket \mathfrak{C}_1, \mathfrak{C}_2, \text{NOT} (\mathfrak{C}_1 \rightarrow \mathfrak{C}_2) \rrbracket.$$

Once again, this is sufficient to *falsify* \mathcal{L} 's claim that S_2 is a *logical consequence* of S_1 . Although in this case S_2 is a material consequence of S_1 , it is neither a *nomical* nor a *logical* consequence of S_1 :

$$(12) \quad (\text{Level of Consequence}) \quad \text{NOT: } S_1 \models_{N,L} S_2.$$

d. Finally, suppose that the world being as S_1 says carries a *nomical* guarantee of its being as S_2 says, but *not a stronger guarantee*—in particular, not a guarantee as strong as that required for a *logical* consequence:²⁴

$$(13) \quad (\text{Level of World}) \quad \llbracket (\mathfrak{C}_1 \rightarrow \mathfrak{C}_2), \text{NOT} (\mathfrak{C}_1 \Rightarrow \mathfrak{C}_2) \rrbracket.^{25}$$

This, too, *falsifies* \mathcal{L} 's claim that S_2 is a *logical* consequence of S_1 . Although S_2 is a *nomical* consequence of S_1 , it is not a *logical* consequence of S_1 :

$$(14) \quad (\text{Level of Consequence}) \quad \text{NOT: } S_1 \models_L S_2.$$

These considerations show that there are a number of ways in which consequence relations are constrained by the world, i.e., in which the world can falsify consequence claims, including claims made by the proposed logical theory, \mathcal{L} . We conclude that for a theory of consequence to be correct, the world has to cooperate. A theory of consequence cannot conflict with those aspects of reality that determine whether given the truth of some sentences, the truth of others is *possible*, is *the case*, is *significantly guaranteed* to be the case, or is *guaranteed with an especially strong modal force* to be the case. That is to say, a correct theory of consequence must respect the connection between the conditions that have to hold in the world for the premise-sentences to be true and those that have to hold for the conclusion-sentence to be true. In this sense, the world *limits* the options open for theories of consequence. In particular, we need to justify our choice of a logical theory by showing that its judgments are not undermined by the world.

Of course, this does not rule out the possibility that there is something in a proposed logical theory, \mathcal{L} , that (i) prevents its judgments from ever being challenged by the world, and (ii) makes all its judgments true with the requisite modal force. But in that case we would have to *veridically justify* the claim that the proposed logical theory is *not, and cannot be, falsified by the world*, i.e., even in this case we have to ground logic in the world on some level in a non-trivial sense.

²⁴ That is to say, reality is such as to guarantee the transmission of truth from S_1 to S_2 in all *physically possible* worlds but not in larger collections of possible worlds.

²⁵ Here " \rightarrow " has the modal force of a nomical law and " \Rightarrow " has a stronger modal force.

Before moving beyond the conclusion that logic is *constrained* by the world, I would like to emphasize again that this conclusion relies on *inherent*, rather than *accidental*, connections between logical consequence and truth and between truth and the world: It is *inherent* in logical consequence (or in its “job description”) that it provides a guarantee of preservation of truth from premises to conclusion and that this guarantee is especially strong; and it is *inherent* in truth that the truth of a sentence substantially depends on the way the world is (or the way relevant aspects of the world are).

B. *Why, and How, Logic is (can be) Enabled by the World.* Having shown that logic is significantly *constrained* by the world, our next task is to show that it is also significantly *enabled* by the world. We need to show that, certain things (like language and meaning) being held fixed, the world can *bring it about* (or play a significant role in bringing it about) that some sentences *follow logically* from some other sentences. This will provide a *positive* explanation of logic’s grounding in reality, and an explanation of how the grounding of logic in reality ensures its *veridicality*.²⁶

Let us begin, once again, with consequence in general, starting with *material* consequence.

Case 1: Material Consequence (\models_M). Let S_1 and S_2 be as above, i.e., their truth value is determined by the way the world is. Suppose that the conditions required for both S_1 and S_2 to be true hold, i.e., the world is as both S_1 and S_2 say it is:

$$(15) \quad (\text{Level of World}) \quad \llbracket \mathfrak{C}_1, \mathfrak{C}_2 \rrbracket.$$

This by itself (or rather largely by itself)²⁷ gives rise to the material consequence:

$$(16) \quad (\text{Level of Consequence}) \quad S_1 \models_M S_2.$$

That is to say, the world can justify claims of material consequence. In other words, the world is capable of “guaranteeing” the transmission of truth from sentences to sentences with a *material* or *factual* force. The world can *bring it about* that some sentences are material consequences of others. For example, the reality of Obama being president and Earth having exactly one moon brings it about that the judgment:

$$(1) \quad \text{Barack Obama is president; therefore, Earth has one moon}$$

is a true judgment of a material consequence, one that holds with a material, or factual, force.

Case 2: Nomic Consequence (\models_N). Let S_1 and S_2 be as above, and suppose that the world is governed by a natural law that connects the conditions required for S_1 to be

²⁶ “Ensure”—not in an absolute, indefeasible sense, as demanded by foundationalists, but in a defeasible, yet a strong, sense, as required by foundational holists.

²⁷ Since the mind plays a relevant role too, e.g., in creating language and connecting it to the world.

true with those required for S_2 to be true.²⁸ That is to say, \mathfrak{C}_1 necessitates \mathfrak{C}_2 with the force of a law of nature. Put differently, the world is governed by a law that positively connects \mathfrak{C}_1 and \mathfrak{C}_2 , a law with the modal force of a natural law:

$$(17) \quad (\text{Level of World}) \quad \llbracket \mathfrak{C}_1 \rightarrow \mathfrak{C}_2 \rrbracket.$$

This (largely) by itself gives rise to the nomic consequence:

$$(18) \quad (\text{Level of Consequence}) \quad S_1 \models_N S_2.$$

So the world can bring it about both that some sentences are *material* consequences of others and that some sentences are *nomic* consequences of others. An example of such a *nomic* consequence (consequence-schema) is:

$$(2) \quad \text{The force exerted by } a \text{ on } b \text{ is } c; \text{ therefore, the force exerted by } b \text{ on } a \text{ is } c.$$

To the extent that Newton's third law holds in the world, the world guarantees that (2) holds with the force of a nomic law.

Case 3: Logical Consequence (\models_L). Let S_1 and S_2 be as above. We have not determined the modal force appropriate for logical consequence yet, but we have already determined that it is stronger than that appropriate for nomic consequence.

So let us provisionally choose some intuitive characterization of logical consequence that reflects this partial determination, one that (in order not to beg the question whether logical consequence is, or is not, grounded in reality) *can*, though does not have to, be understood in terms that pertain to reality. Say, logical consequence is formally and universally necessary. Let us call this characterization X. Now, suppose that there are *laws in the world* that satisfy X. Then, in principle, these laws can connect the conditions that must hold in the world for some sentences to be true with those required for other sentences to be true. In the case of S_1 and S_2 , suppose there is an X-law that connects the conditions that must hold for S_1 to be true to those that must hold for S_2 to be true:

$$(19) \quad (\text{Level of World}) \quad \llbracket \mathfrak{C}_1 \rightarrow \mathfrak{C}_2 \rrbracket.$$

I.e., the world is governed by a law that positively connects \mathfrak{C}_1 and \mathfrak{C}_2 with a stronger modal force than that of laws of nature, a force sufficiently strong for logic to play its designated role in knowledge. Then this law, (largely) by itself, will give rise to a consequence whose force is that of a logical consequence:

$$(20) \quad (\text{Level of Consequence}) \quad S_1 \models_L S_2.$$

It follows that if, and to the extent that, the world is governed by appropriately strong laws, it can, and does, give rise not just to *material* and *nomic* consequences but also

²⁸ In "possible worlds" terminology, suppose the two conditions are connected by a law that holds in all physically possible worlds or domains.

to *logical* consequences. That is to say: in principle, a sufficiently strong law governing the world can substantiate consequence claims with the force appropriate for logic. Are there such laws? At least on the intuitive level it appears that there are. Consider an example. Suppose the world is governed by the law:

$$(21) \quad (\text{Level of World}) \quad \text{Non-empty } (B \cap C) \rightarrow \text{Non-empty } (C),$$

where B and C are any first-level properties and \cap is the operation of meet on properties, viewed extensionally. This intuitively formal law²⁹ is naturally viewed as universal and as having an especially strong modal force. Therefore, intuitively, this law guarantees the truth of:

$$(3) \quad (\exists x)(Px \ \& \ Qx); \text{ therefore, } (\exists x)Qx,$$

which is one of our intuitive paradigms of logical consequence.³⁰ More generally, this law, and others like it, appear to have an especially strong modal force, one appropriate for a theory that plays the role logic is designed to play in our system of knowledge.³¹

We have seen that, and how, the reality of certain *laws* or regularities *in the world* has the power to establish at least some paradigmatic claims of logical consequence, ones that accord with logic's designated role in our system of knowledge. What kind of laws these are is something we will investigate in the following sections, and the results of these investigations will direct us to observations (made earlier in this essay) that support the existence of such laws. For the time being we can conclude that theoretically, logic is grounded in reality at least in the negative sense of being constrained by reality and, subject to further investigations (which we will turn to shortly), it is also grounded in reality in the positive sense that certain laws governing reality sanction, or give rise to, logical consequences. We have already seen that to fulfill its designated role logic must be *in sync* with the world. We now see that being *in sync* with the world (with certain especially strong laws governing the world) will *enable* logic (or play a central role in enabling it) to perform this role.

All this is perfectly compatible with Thesis 2, which says that logic is grounded not just in the world but also in the mind. Logic is grounded in the world in a way that is shaped by the structure of our mind, and it is grounded in the mind in a way that is both constrained by and takes advantage of certain structural features of the world.

²⁹ At this point, we can understand "formal" as in Chapter 8, Section 4. Later on I will offer a precise characterization of formality.

³⁰ For the sake of simplicity I talk about schematic claims as true/false.

³¹ Note: someone might argue that this analysis is read off the standard (Tarskian) model theory of logic, and it is true that one can naturally connect the two. But to treat this analysis as read off this semantics is to put the cart before the horse. Rather, it is this analysis, or the systematic foundational account it leads to, that *explains* why and how model-theoretic semantics offers a sound understanding of logic. (For more on the philosophical significance of Tarskian models see the next section.)

Still, someone might raise the following objection: You have shown that logic must take the world into account and that the world could, in principle, ground it both negatively and positively. But you have not ruled out the possibility that while the world is capable of grounding logic, this job is *in fact* done by the mind alone, with the world playing no role. Suppose, for example, that nature endowed us with a built-in logical apparatus that happened to work in the world. In that case we would not need the world to either sanction or constrain our logic. Logic would not conflict with the world, but this would be in virtue of the structure of our mind. Similarly, although the world could positively ground logic, it would be the structure of our mind that did the job.

My response to this objection is: whether we have a built-in logical system that works properly in the world or we have to build one for ourselves by consulting the world, the *veridicality* of such a system significantly *depends* on the world. Were the world different in relevant ways, were it governed by different laws, the logic generated by our mind would fail in it. Even if evolution or a benevolent God saw to it that if reality changed, our mind changed with it, it would still be *reality* (and not just our mind or evolution or God) that would ground our logic. It is only if God, evolution, or our mind could, and would, change the world so that it always obeyed our built-in logic, that the world would not play a (central) role in grounding logic. And even then it would, since after all it would be the ability to *regulate the world* that was responsible for the adequacy of logic.³²

We thus confirm our conclusion that claims of logical consequence must ultimately be judged by reference to the world: for S_2 to follow logically from S_1 , the conditions that have to hold in the world for S_1 to be true cannot conflict with the conditions that have to hold in the world for S_2 to be true. No matter how our logical theory is generated, this condition must be satisfied. This is one of the major *friction* requirements on logic. And it is this requirement that explains, from the perspective of logic itself (and not just from a general epistemic perspective), why logic cannot be confined to the *center* but must lie in the *periphery* as well.

The need to ground logic in reality also explains why logic *is not*, and *cannot be*, *analytic*, at least in the traditional sense in which analyticity rules out any constraints by the world. In a recent exchange in the Foundations of Mathematics (FOM)³³ it was remarked that if arithmetic is part of logic then “[t]he question of ‘justifying the axioms’ . . . goes away . . . because **logic comes ‘for free’**”. This is a view commonly associated with the claim that logic is analytic, and it is just this view that our considerations reject. For logic to *work in the world*, for logic to *avoid factual errors*,

³² Another way to put this is to say that even if a *genetic* account of logic (an explanatory, historical-like account of how logic was developed or learned) might focus on the mind, a *justificatory* account of logic must significantly appeal to the world. (This is related to the genetic fallacy argument discussed in Chapter 3, Section 4.)

³³ An email list for discussing the foundations of mathematics, moderated by Martin Davis.

for logic to *transmit correspondence-truth with a strong modal force*—in short, for logic to play its designated role in knowledge—it cannot come *for free*. Logic, like any other *veridical* discipline, must face considerable *factual* constraints, constraints coming from the *world* and not just from mind or language. This reason for rejecting the view that logic is analytic joins other, more general, reasons for rejecting analyticity as an epistemically significant category (see Part II). Together, they speak against the traditional view that logic is grounded exclusively in the mind (language, convention, brain structures) and in favor of the non-traditional view that logic has, and must have, a significant grounding in reality.

To reject the claim that logic is analytic, however, is not to deny its significant linguistic orientation. Logic is a theory of the logical properties of, and the logical relations between, linguistic entities, and in this sense its subject matter is linguistic. This might tempt us to think that logic's veridicality is grounded in laws governing language. But a deeper understanding shows that language is not the only, or even the main, source of its veridicality.

Our next task is to find out what features of reality logic is grounded in, and how. In pursuing this task, we will be interested in features that ground logic not just negatively, but also positively.

10.4 What Specific Features of the World is Logic Grounded in?—The Formality Thesis

Our considerations in the last section have shown that an appropriate foundation for logic would ground it in universal and modally strong objectual laws, laws connecting the circumstances that have to hold in the world for the premises of a given logical consequence to be true to the circumstances that have to hold in the world for its conclusion to be true. The question we now face is: what objectual laws would fit the bill? Our proposed answer is:

Thesis 3: Logic is Grounded in Formal Laws

Logic (logical consequence) is grounded in formal laws governing reality—laws governing formal features of objects and properties—and it is the broad applicability and strong modal force of such laws that underlies (and explains) the generality and strong modal force of logic.

We call this the “Formality Thesis”. To render the formality thesis an informative and systematic thesis, we have to address a number of questions. What is formality, theoretically? Why are formal laws highly general? Why do they have an especially strong modal force? How do we build a logical system based on formal laws? Is our current paradigm of a logical system built in this way? Does it exhaust the array of valid inferences based on formal laws? And so on. We will devote the present section and some of the sections following it to answering these questions, but first let us remind ourselves what we have said about the formal earlier in this essay and provide a few examples.

In developing our model of knowledge in Part II, we said (Chapter 3) that each theory in our corpus of knowledge must square off with some facet of reality: *logic* with the *formal behavior of objects*, physics with their physical behavior. We explained how *logic* can be vulnerable to the world by saying that if the *formal behavior of objects* conflicts (in some “deep” way) with a *particular logic*, then *this logic* is challenged by the *behavior of those objects*. And we explained how logic lies both in the periphery and in the center of our model of knowledge by saying that as a theory of the *formal laws governing structures of objects*—for example, the laws relating identity to cardinality—logic is in the *periphery*; as a theory of the logical laws governing statements and sets of statements, logic is in the *center*. As a theory saying that a structure whose objects are all identical is a structure of one object, logic is about the world; as a theory saying that the statement “All objects are identical” implies the statement “There is exactly one object”, logic is about language. As a theory of the *formal structure* of “real” objects (configurations of objects), logic is anchored in *reality*; as a theory of the formal structure of our thought of objects, logic is anchored in the mind.

In discussing the conception of reality characteristic of our model (Chapter 5, Section 2), we asked whether there were strong commonsensical reasons for presuming that reality—including everyday physical reality—had *abstract features*, and we took *formal features* as a *paradigm* of such features. Our answer was positive. We said that it did not take more than plain common sense to see that physical objects and their physical properties had formal features. To demonstrate this we took as our starting point non-controversial individuals: students in a certain graduate seminar. We observed that each student in that seminar had properties of many kinds, including formal properties, and that many of the properties of students in that seminar, including most of their physical properties, had properties of a variety of kinds as well, including formal properties. For example, each individual in the class had the *formal property* of self-identity. The property of being a student in the class had the *formal cardinality* property SEVENTEEN (there were seventeen students in the seminar). The relation of studying-in-the-same-class-as had the *formal* properties of being reflexive, symmetric, and non-transitive, and so on. Similarly, properties of students in the class could be operated on, or be combined, by a variety of *formal operations*: complementation, union, intersection (or more generally, formation of Cartesian products), etc. For example, by applying the *formal* operation of intersection to the properties of being a first-year student and being a woman we obtained the intersective property of being a first-year female student. We concluded that if the students in the seminar in question were real and if they had the (non-formal) properties and relations mentioned above, then it was hard to contest that they, and their properties, had formal properties as well, and that these formal properties were real. These formal properties were just as factual as all the other properties involved in that part of reality which involved this class.

Next, in the section on truth in mathematics (Chapter 8, Section 4) we observed that formal features, like physical features, had the potential of exhibiting regularities

and being governed by laws, and that this potential was in fact realized: self-identity is a law, cardinalities are governed by various laws of order, and so on. Furthermore, if formal features were governed by regularities and laws, then it was sensible to presume that these regularities and laws, like many other regularities and all laws, had a certain modal force, a modal force that went beyond their application to actual objects.

These observations form the intuitive basis for our conception of formal features and laws, and for our belief that such features and laws are real and as such available for a foundation for logic (that is, a holistic foundation).

I should note in passing that the fact that both mathematics and logic, on our view, are grounded in formal laws is philosophically significant. It opens new possibilities for our understanding of the relation between logic and mathematics, something I will exploit in Section 10.8.

Focusing our attention on logic, our next task is to provide a few intuitive examples of formal laws (putatively) grounding logical consequences. In the last section we came upon one example of such a law:

$$(21) \text{ Non-empty } (B \cap C) \Rightarrow \text{Non-empty } (C).$$

This law (putatively) grounds the logical consequence:

$$(3) (\exists x)(Px \ \& \ Qx) \models_L (\exists x)Qx.$$

Similar laws ground other logical consequences. For example,

$$(22) \text{ Non-empty } (B) \Rightarrow \text{Non-empty } (B \cup C),$$

and

$$(23) B \subseteq C, \text{ Universal-in-A } (B) \Rightarrow \text{Universal-in-A } (C),$$

where A is an arbitrary domain (non-empty collection) of individuals and B, C are arbitrary properties with extensions restricted to A ,³⁴ ground the logical consequences

$$(24) (\exists x)Px \models_L (\exists x)(Px \vee Qx),$$

and

$$(25) \{(\forall x)(Px \supset Qx), (\forall x)Px\} \models_L (\forall x)Qx,$$

respectively.

So far we have considered consequences whose premises could all be true. What about

$$(26) Pc \ \& \ \sim Pc \models_L S,$$

³⁴ Here I am interpreting the universal quantifier as associated with a domain (universe). I will return to this issue later in the chapter.

where S is an arbitrary sentence? This consequence could be viewed as being grounded in the *formal impossibility* of an individual having both a property and its complement in any (non-empty) domain.³⁵

We can also explain *failures* of logical consequence by referring to formal laws. For example, the claim

$$(27) \quad (\exists x)(Px \vee Qx) \models_L (\exists x)Px$$

is false because there is no objectual formal law that grounds it. In particular:

$$(28) \quad \text{Non-empty } (B \cup C) \Rightarrow \text{Non-empty } (B)$$

is not a formal law. The same holds for the consequences

$$(29) \quad (\exists x)Px \models_L (\exists x)(Px \& Qx)$$

and

$$(30) \quad (\exists x)Px \models_L Pc.$$
³⁶

Having seen a few intuitive examples of the formality thesis at work, we proceed, next, to a *theoretical* characterization of the formal. The formality thesis says that logic is grounded in formal laws governing reality, laws governing the formal features of objects and properties in the world. But what, theoretically, do we understand by “formal”?

In setting out to provide a theoretical characterization of formality, however, we face a methodological dilemma. On the one hand, we aim at a general theoretical characterization, one that is not committed to any specific mathematical theory of formal structure. On the other hand, we aim at a precise characterization, one that will identify the scope of formality and explain its nature in exact terms. But the

³⁵ To simplify the presentation I limit myself here and below to classical examples, but this is not essential for the account.

³⁶ We can explain the logical difference between

$$(31) \quad \text{Some } \textit{fake} \text{ roses are fragrant} \models_L \text{Some roses are fragrant}$$

and

$$(32) \quad \text{Some } \textit{red} \text{ roses are fragrant} \models_L \text{Some roses are fragrant}$$

in a similar way. Due to the fact that a fake rose is a non-rose and that “non”, in the above sentence, refers to a formal operator (complement in a given universe), the logical structure of (31) is

$$(33) \quad (\exists x)(\sim Rx \& Fx) \models_L (\exists x)(Rx \& Fx),$$

in contrast to the structure of (32), which is

$$(34) \quad (\exists x)(Dx \& Rx \& Fx) \models_L (\exists x)(Rx \& Fx).$$

The falsehood of (31) is due to the fact that the conditions that must hold in the world for its premise to be true do not formally necessitate the conditions that must hold in the world for its conclusion to be true. In particular:

$$(35) \quad \text{Non-empty } [(\text{Complement-in-A of } B) \cap C] \Rightarrow \text{Non-empty } (B \cap C)$$

is *not* a formal law.

second aim requires the use of sophisticated resources, and this means using a specific, well-developed mathematical theory of formal structure as our background theory, when in fact we are seeking an account that leaves it an open question which mathematical theory of formal structure is correct or is overall the best theory available today.

One serious problem with presenting our characterization in terms of a particular theory of formal structure is that such a theory, like any specific theory of any complex subject matter, is likely to have certain weaknesses and deficiencies. Clearly, it is a mistake to reject a given general, philosophical conception of logic on the ground that a particular theory used to precisify it has certain peculiarities. But it is tempting to do so. For example, it is tempting to reach a sweeping negative conclusion about the viability of a conception of logic that appeals to modal force on the ground that a given background theory—say, Zermelo-Fraenkel set theory with the axiom of choice or ZFC—(purportedly) fails to provide adequate resources for dealing with modal force.³⁷ But a sweeping conclusion of this kind is likely to be unwarranted: one could, in principle, use, or develop, a different mathematical background theory, or revise ZFC so that it no longer has this weakness. More fundamentally, it is a universal predicament of human theorizing that hardly any of our background theories are perfect, yet this is not a reason for discarding all our theories or drawing sweeping negative conclusions about their viability.

To avoid this problem, I will present my theoretical characterization of formality on two distinct levels: a general philosophical level, and a specific mathematical level. First, I will offer a general, non-technical account of formality as a basis for logicality and identify some of its philosophical roots. Next, I will use the resources of a sophisticated (though imperfect) mathematical theory—classical set theory, or more specifically, ZFC—to formulate a more precise account and show how questions that are left open by the general account can be resolved. This two-step account will enable those who reject, or find flaws in, ZFC to understand our formality principle *in general terms* and to use our set-theoretic formulation as an *example* of how to construct a precise version of this principle, employing the resources of whatever they take to be a correct theory of formal structure. Just as importantly, this will enable us to highlight the philosophical nature of our foundation and to put the relation between the general principles and their instantiations in proper perspective.³⁸ In pursuing this two-tier account, I will proceed with the general account as far as possible, then turn to the specialized account, and end with questions involving both.³⁹

³⁷ We will discuss a claim like this in Section 10.7. See also the discussion of other purported problems with ZFC in Section 10.7.

³⁸ To simplify the discussion, I will assume bivalence, but it will always be possible to reformulate the account so that neutrality with respect to bivalence is achieved.

³⁹ Before or while reading this subsection and the following one, the reader may wish to consult the informal characterization of terms (“object”, “feature”, “property”, “level”, etc.) we gave in our discussion of mathematics in Chapter 8, Section 4.

1. *The formality thesis: a general theoretical account*

1. *What is Formality, Theoretically?* Our first task is to explain, in general yet informative terms, what a formal property or operator⁴⁰ is.

Formality. Roughly, our proposal is that a *property is formal iff whether an object satisfies it depends only on its structural features, where this is a highly restrictive condition*. To give informative content to this characterization, we will use a notion known for its informativeness, wide applicability, and explanatory power—“invariance”. The underlying idea is that properties in general are *selective* in what they “pay attention to”. For example, the property of being a student pays attention to the educational status of a given person but not to her hair color, height, age, nationality, etc. This enables us to characterize different types of properties according to what features of objects, or what changes in objects, they are, or are not, attuned to. In the case of formal properties, we can characterize them by saying that they are not attuned to, or are *invariant* under, any changes in objects besides the relevant structural changes. The next step is to characterize “invariance under non-formal changes” in a way that explains what the relevant characteristic of a formal structure is.

The defining characteristic of a formal structure is, we propose, *indifference to the distinctive features, or identity, of individuals*, and this can be expressed by saying that a *formal structure is indifferent to 1-1 replacements of individuals*. If you have a formal structure of people and you replace the people, in a 1-1 manner, with, say, numbers, the formal structure will not change. Accordingly, a formal property will not “notice” such a change. A property of people (or of sets, properties, relations of people) is *formal* iff it is preserved under any 1-1 replacement of people by individuals of any kind (say, numbers). In other words, as far as formal properties and operators are concerned, there is no difference between an individual person and an individual number. Formal operators distinguish *patterns* delineated by individuals and their properties, but not the individuals themselves. If we replace any argument of a formal operator by any other which is its image under some 1-1 replacement of individuals by individuals (e.g., a 1-1 replacement of humans by numbers), the formal operator will assign the same value to the two arguments.

Our general characterization of formality is, then: To be formal is not to distinguish any 1-1 replacement of individuals or any replacement of sets, properties, relations, functions induced by such a 1-1 replacement. In short:

(FORMALITY) *An operator is formal iff it is invariant under all 1-1 replacements of individuals.*

⁴⁰ We will move freely between talk of properties and talk of operators. Operators are objectual functions, including functions representing properties, relations, and possibly other types of objects. For example, an operator representing a property *P* is (assuming bivalence) a function that assigns to a given object the value *T* if it has this property and *F* otherwise. Thus, an operator representing the first-level property *is a student* assigns to a given individual the value *T* iff it is a student, and an operator representing the second-level property *is not empty* assigns to a given first-level property the value *T* iff it is not empty.

One example of a formal operator under this characterization is the second-level operator of *non-emptiness*, i.e., the operator denoted by the existential quantifier of standard first-order logic. This operator, whose arguments are first-level properties (extensions of 1-place first-order predicates), is formal because it does not distinguish between two first-level properties whose extensions can be obtained from each other by a 1-1 replacement of their members. If one is non-empty, so will the other be. The second-level operator *is-a-property-of-humans*, on the other hand, is not formal. Consider a domain, D1, consisting of ten men and ten women, and replace its (human) members by numbers (numerical individuals) in a 1-1 manner so that each man is replaced by an even number and each woman by an odd number. Call the image of D1 under this replacement D2. Now consider the first-level property *is-a-man*. This property is assigned the value True by the second-level operator *is-a-property-of-humans* in D1, but its image under the 1-1 replacement described above is assigned the value False by this operator. The image of the first-level property *is-a-man* (restricted to D1) is the first-level property *is-even* (restricted to D2), and the latter is not a property of humans. The operator *is-a-property-of-humans*, therefore, is not formal.

The first-level operator of *identity* and the second-level operator of *intersection*, in contrast, are formal: individuals stand in the relation of identity iff their counterparts under any 1-1 replacement stand in this relation. Likewise, the image of an intersection (of properties) under a 1-1 replacement of individuals is an intersection (of properties, namely, an intersection of the images of the former properties under the given 1-1 replacement of individuals). But the first-level operator *is-taller-than* and the second-level operator *is-a-biological-property* are not formal.

Logicity. Turning to logicity, my proposal is that logic is attuned only to formal features of reality, i.e., logic is indifferent to the identity of individuals and their “personal” features, those features that are tied up with their identity. In the case of logical consequence, logic pays attention only to the *formal skeleton* of the situations that have to be the case for the sentences involved to be true (false). For example, if the sentence “John is a male student” appears in an inference, logic pays attention only to the fact that its truth conditions involve an individual standing in the intersection of two properties of individuals. This is the formal skeleton of its truth conditions, the formal skeleton of what has to be the case for this sentence to be true.⁴¹

The view that logic is indifferent to the identity of individuals (and their “personal” features) has long roots in the philosophical and logical literature. A few examples are:

[G]eneral . . . logic . . . treats of understanding **without any regard to difference in the objects** to which the understanding may be directed (Kant 1781/7: A52/B76).

⁴¹ Logic might also be interested in whether the same individual stands in two (distinct) intersections of properties, but not in “who” this individual is.

[P]ure logic...disregard[s] the particular characteristics of objects (Frege 1879: 5).

[The relation of **logical consequence** between a sentence X and a class of sentences K] **cannot be influenced in any way by... knowledge of the objects to which the sentence X or the sentences of the class K refer.** The consequence relation cannot be affected by replacing the designations of the objects referred to in these sentences by the designations of any other objects (Tarski 1936a: 414–5).

[Logical] quantifiers should not allow us to distinguish between different elements of [the universe] (Mostowski 1957: 13).

Whatever else logic is for these authors, the idea that logic abstracts from the identity of individuals is part of their legacy.⁴² My current proposal is that this trait is the defining characteristic of that aspect of *formality* which is fundamental to *logic*.⁴³ Speaking about operators, I will characterize *logicality* as follows:

(LOGICALITY) *An operator is an admissible logical operator iff it is formal.*

So far we have approached logic as if it were essentially “predicate logic” rather than “sentential logic”. To extend our account to sentential logic, we observe, first, that sentential and predicative logic have two distinct conceptions of language and the world, or operate on two distinct levels. Predicate logic views the world in terms of individuals and their properties (relations), while sentential logic views it in terms of facts or states of affairs. This affects the shape the formality principle takes in the two logics.

In both logics, logical operators are formal in the sense of taking into account only the *formal* pattern of their arguments, where this is characterized as invariance under replacements of atomic elements, or indifference to the identity of such elements. But in predicate logic the atomic elements are individuals, while in sentential logic they are atomic states of affairs. Accordingly, we characterize formality for *predicate logic* by saying that formal operators take into account only the pattern of objects having properties and standing in relations, disregarding “who” the individuals involved are. We characterize formality for *sentential logic* by saying that formal operators take into account only the pattern of states of affairs being, or not being, the case, not their identifying features (those distinguishing one atomic state of affairs from another).⁴⁴

⁴² In his dissertation MacFarlane distinguishes three senses in which logic has historically been said to be formal: *Formality 1*—Logic is formal “in the sense that it provides constitutive norms for thought as such”. *Formality 2*—Logic is formal “in the sense that it is indifferent to the particular identities of objects”. And *Formality 3*—Logic is formal “in the sense that it abstracts entirely from the *semantic content* of thought” (MacFarlane 2000: ii). The sense of formality we are concerned with here is closest to MacFarlane’s Formality 2. For another historical discussion of “logical formality” see Dutilh Novaes (2011).

⁴³ For a more general theory of formality that coincides with ours with regard to the formality of logic, see Beck (2011).

⁴⁴ Notes:

- (i) As before, I will leave ontological questions—in particular, ontological questions concerning individuals, properties, and states of affairs—for a later discussion, where the issue of epistemic freedom (essential to understanding *representation*) will play a more central role.
- (ii) What is treated as an *atomic* state of affairs (like what is treated as an *individual*) is relative to context.

We thus say that a *sentential operator* is *formal* iff it is invariant under all 1-1 replacements of atomic states of affairs that preserve being/not-being the case (abbreviated as $+/-$). This characterization coincides with the standard Boolean criterion for sentential logical operators, which codifies the notion of *formality* for such operators.

Putting together the formality characterizations of sentential and predicative operators and using “+” and “-” for “is the case” and “is not the case”, we can capture the similarity between them by:

- (FORMALITY*) (a) A sentential operator is formal iff it is invariant under all replacements of atomic states of affairs that preserve $+, -$.
 (b) A predicative operator is formal iff it is invariant under all replacements of individuals that preserve identities and non-identities.

Although predicative and sentential operators stand on a par, we will continue to pay more attention to logicity and formality as they apply to the former. The reason is that both from our perspective on truth and knowledge and from the perspective of logic itself, the conceptual scheme of objects and properties has priority over that of facts or states of affairs. From a logical point of view, this is one of the lessons we learned from Frege and Russell: it is only by dealing with objects and their properties (in a broad sense that includes relations) that logic achieves full fruition.⁴⁵

2. *How to Build a Logical System Based on Formal Laws.* Our next task is to describe the structure of a logical system based on the formality thesis in theoretical, yet fairly general, terms. This we do by showing how the standard template of a logical system—syntax, semantics, and proof method—should be configured in order to represent the idea of a method of logical inference grounded in formal laws. In describing the structure of such a system I will not assume a particular background theory (like ZFC).⁴⁶

SYNTAX

The syntax is designed to identify the *logical structure* of sentences, where this structure reflects the *formal structure* of the situations that have to hold in the world for different sentences to be true/false. (More generally, logical structure

⁴⁵ The priority of predicative to sentential operators finds several expressions in logic itself. For example, every logical sentential operator corresponds to a logical predicative operators, but not vice versa: “ \sim ” corresponds to *complementation* and “ $\&$ ” to *intersection* (in “ $Px \& Qx$ ”), or more generally *Cartesian product* (in “ $Px \& Qy$ ”), but “ $=$ ” and “ \exists ” do not correspond to any sentential operators (on finitely long sentences). This is related to the fact that the definition of truth for predicative languages covers sentential sentences but not the other way around. We can say that “ $\sim Pa$ ” is true iff the referent of “ a ” is in the complement of the extension of “ Px ” in a given universe. But we have no definition of truth in (finitist) *sentential* terms for sentences containing predicative operators.

⁴⁶ However, to save time and space I will rely on readers’ familiarity with standard textbook examples of logical systems (which often assume such a background theory) and on their ability to identify the *general structure* of such systems (which is not dependent on the specific details of such a theory).

reflects the *formal skeleton* of the truth conditions of sentences of a given language.) To this end we divide the constants of the language into *logical* and *non-logical*, building the former as “*fixed*” constants of the system and the latter as “*non-fixed*”. Logical constants denote the *distinguished*—i.e., *formal*—*parameters* of the situations that have to hold in the world for the sentences involved to be true/false, while non-logical constants denote the non-distinguished parameters of these situations.⁴⁷ Accordingly we characterize “admissible” logical constants based on their denotations as follows:

(LOGICAL CONSTANT) *A constant is an admissible logical constant iff it refers to (denotes, represents) a logical—i.e., formal—operator.*

Different logical systems can include different sets of logical constants, and together these systems form what we may call “a family” of admissible logical systems. Further constraints on logical constants have to do with their ability to function properly in the semantic and proof-theoretic subsystems, and we will specify some of these below.

SEMANTICS

The semantics and proof theory provide two methods for identifying logical consequences (logically valid inferences)—the former in terms of *truth* and the latter in terms of *proof*. To the extent that we have identified the role of logic in terms of truth, the semantic approach is integral to our account. (But this does not detract from the importance we attach to the proof method, as we will see shortly.) One kind of semantic subsystem that enables logic to play its designated role in knowledge has the following components:

- (a) An apparatus of models.
- (b) Constraints on logical constants.
- (c) A definition of “truth in a model”.
- (d) A definition of “logical consequence”.

- (a) *Models*. The apparatus of models is so constructed as to represent the totality of *formally possible* states of affairs vis-a-vis the language of the system, where this totality itself represents the full scope of formal laws, actual and counterfactual, relevant to the logical system in question.⁴⁸
- (b) *Constraints on Logical Constants*. Logical constants are subject to constraints concerning their proper functioning in the semantic subsystem:

⁴⁷ Although in exceptional cases, such as “ Pa ; therefore Pa ”, logical consequences do not involve logical constants, normally they do.

⁴⁸ For the sake of brevity, I will often use “counterfactual” in a way that includes “actual”.

- (LOGICAL CONSTANT*) *A logical constant C must satisfy the following conditions:*
- (i) *C denotes a formal operator.*
 - (ii) *C is extensional.*
 - (iii) *C is defined over all models.*
 - (iv) *C is assigned a “pre-fixed” denotation in all models.*
 - (v) *C is a rigid designator.*

Conditions (ii) through (v) are explained in Sher (1991, Chapter 3, 2001, and 2003). We will discuss some misunderstandings concerning these conditions in Section 10.7.

(c) *Truth in a Model.* To identify logical consequences semantically, in terms of models, we need a definition of “truth in a model”. A natural way to arrive at such a definition is to take a definition of truth (simpliciter) which is *tuned to the logical structure of sentences*, or more specifically, to the way logical structure affects the truth value of sentences (as in Tarski’s definition of truth, on our interpretation)⁴⁹, and relativize it to models. In so doing, we make sure that logical constants are assigned their prefixed denotations in all models.⁵⁰

(d) *Logical Consequence.* The definition of logical consequence is the usual Tarskian definition:

- (LOGICAL CONSEQUENCE) *σ is a logical consequence of Σ iff there is no model in which all the sentences of Σ are true and σ is false.*

Under this definition (given (a) through (c) above), logical consequence is grounded in laws (regularities) that hold in all formally possible situations, and as such it is strongly necessary, as the job description of logic requires.

Putting syntax and semantics together, we can describe the resulting subsystem as follows: The syntactic-semantic subsystem has two parts, each of which is itself a complex, syntactic-semantic structure. One structure consists of logical constants (syntax) and an apparatus of models (semantics). The second structure consists of non-logical vocabulary (syntax) and the denotations of non-logical constants within models (semantics). Syntactically, logical constants are formula-building operators; semantically, they are assigned *pre-fixed* formal operators on models which represent formal properties of elements of models. Since the logical constants denote formal properties of the denotations of the non-logical constants, their level is usually higher than that of the non-logical constants (serving as their arguments).

The syntactic-semantic subsystem is generated by combining these two structures. Syntactically, this is done by rules for forming well-formed formulas by means of the logical operators, and semantically, by rules for determining truth in a model based

⁴⁹ See Chapter 8, Section 5.

⁵⁰ There is no need to go into details, since this procedure is implemented in standard textbooks of mathematical logic.

on the formal denotations of the logical vocabulary. To satisfy the requirement that logical consequences have a strong modal force, it is essential that the denotations of the extra-logical vocabulary vary from model to model, i.e., that non-logical constants be “strongly variable”. It is also essential that the denotations of both logical and non-logical constants be defined over models, and indeed all models, so that every formally possible state of affairs will be taken into account in determining logical truths and consequences. To satisfy the formality condition it is essential (i) that the semantic values of the logical and non-logical constants be extensional, (ii) that logical constants be assigned formal denotations in each model, (iii) that the denotations of the logical constants be determined in advance for all models (rather than separately within each model), and (iv) that they be fixed throughout the totality of models (rather than varying arbitrarily from model to model). We may say that logical constants are semantically defined *outside* models, on the same level as the definition of models, and that non-logical constants are semantically defined *within* models, hence on a different level from the definition of logical constants.

Now, if we use “ ℓ_1, \dots, ℓ_n ” to represent the logical parameters (constants) of a sentence σ and a set of sentences Σ , “ \mathfrak{C}_1 ” and “ \mathfrak{C}_2 ” to represent the situations that have to hold in the world for σ and the sentences of Σ to be true, “ f_1, \dots, f_n ” to represent the formal parameters of \mathfrak{C}_1 and \mathfrak{C}_2 (corresponding to ℓ_1, \dots, ℓ_n), and “--- [...]” to represent the relevant ordering of the logical parameters of Σ and σ , then we can represent the grounding of logical consequences in formal laws figuratively as follows:

(36) (<i>Level of Language</i>)	$\Sigma[\ell_1, \dots, \ell_n] \models_{\perp} \sigma[\ell_1, \dots, \ell_n]$
	\Updownarrow
(<i>Level of Models</i>)	Truth is preserved from $(\Sigma[\ell_1, \dots, \ell_n])$ to $\sigma[\ell_1, \dots, \ell_n]$ in all models (representations of all formally possible counterfactual situations relevant to Σ and σ)
	\Updownarrow
(<i>Level of World</i>)	There is a formal law which, given $\mathfrak{C}_1[f_1, \dots, f_n]$, necessitates (with the modal force of a formal law) $\mathfrak{C}_2[f_1, \dots, f_n]$.

The main ideas embodied in (36) are: (i) logical connections between sentences (relevant to logical consequence) are due to their logical structure, (ii) formal connections between situations are due to their formal structure, (iii) formal connections that hold in all models are based on formal laws, (iv) such laws have a generality and modal force proportional to their scope of applications, and (v) the relation of logical consequence between σ and Σ is grounded in a formal law that connects the formal structures of the conditions that have to be satisfied in the world for all the sentences of Σ to be true and those that have to be satisfied for σ to be true. Since this law is formal, it connects the formal structures in question with considerable generality and modal force, as required by an adequate grounding of logic.

The semantic definitions of logical truth, logical contradiction, logical inconsistency, etc. are derivative from the semantic definition of logical consequence:

(LOGICAL TRUTH) σ is logically true iff σ is true in every model.

(LOGICAL CONTRADICTION [FALSEHOOD]) σ is logically false iff σ is false in every model.

(LOGICAL INCONSISTENCY) Σ is logically inconsistent iff there is no model in which all the sentences of Σ are true.

These definitions express the principles that logical truths are grounded in formal laws and logical contradictions/inconsistencies are grounded in formal impossibilities.

PROOF THEORY

The proof-theoretical subsystem is designed to provide a decidable set of axioms and rules of derivation that yield finitely checkable proofs of logical consequences (logical truths, contradictions, and inconsistencies). The grounding of the proof-theoretic subsystem in reality is philosophically straightforward: logical axioms and rules of derivation express (encode, represent) formal laws. This connection is codified by the metalogical requirement of *soundness*: to be sound, the proof method must license only semantically valid consequences, that is, on our account, consequences grounded in formal laws, and as such truth-preserving with an especially strong modal force.

In the philosophical literature the proof method is rarely viewed as grounded in formal laws governing reality. But viewing it in this way has considerable advantages, not just from the perspective of an overall philosophical foundation for logic, but also from the perspective of understanding proof theory itself: (a) Explaining the proof method by appealing to laws governing reality spares us from the need to appeal to obviousness, convention, exclusively pragmatic considerations, unexplained “intuition”, etc. These, as we have seen in Chapter 9, Section 2, are all highly problematic, and some of them neglect important constraints on logical inferences (consequences), especially constraints having to do with veridicality. (b) Viewing the proof method in this way gives a deeper, and more critical, sense to its *formality* than does the common, purely syntactic account. The proof method is formal not simply in the sense of dealing with *strings of symbols*; it is formal in manipulating symbols that represent *formal* features of objects, and it manipulates them based on *formal* laws—laws that govern these formal features. (c) Viewing the proof method as grounded in the same laws that the semantic method is grounded in unifies the semantic and proof-theoretic subsystems of logic, rather than pitting them one against the other. And it does this in a way that recognizes the centrality of both, giving preference to neither. Each system is grounded in the formal in its own, equally basic, way: the proof method by *encoding* formal laws, the semantic method by *representing* formal laws as *regularities in all models*.

II. The formality thesis: a specialized mathematical perspective

Our general account of logic as grounded in formal laws did not appeal to any highly specialized mathematical theory in characterizing formality. It did exploit our general

understanding of basic mathematical concepts, but in a way that did not depend on which highly specialized mathematical theory (or theories) they belong to. In developing an actual logical system based on the principles delineated in the general account, however, one has to use a specialized background theory to render the system precise, fill in details that were not part of the general account, and answer specific questions that are left open by that account. The general account constrains our choice of such a background theory but is in principle compatible with a wide array of such choices. Inevitably, the background theory we choose will not be perfect, and this means that the resulting logical system will not be perfect. But, being human, what we can aim at in constructing an actual logical system is that the general principles it is based on be sound, that it instantiate these principles in a way that is roughly adequate, and that it have the potential of constantly improving through, among other things, revision, or replacement, of its current background theory of formal structure.

One candidate for a background theory of formal structure is ZFC. Its adequacy is, and will always remain, an open question. But it has several virtues that render it a natural candidate for this task in our time. These include its ability to account for classical mathematics in its entirety (Gödelian incompleteness notwithstanding), its proven record as a background theory of the syntax, semantics, and proof theory of standard logic, its fruitfulness in fields of knowledge outside logic (e.g., linguistics), its success in dealing with the well-known paradoxes that threatened logic in the past, the fact that no equally disastrous paradoxes affecting ZFC itself have been discovered, the existence of relative-consistency proofs for its axioms, and so on.

Now, in describing the development of an actual logical system, or a family of such systems, using ZFC as a background theory, we may distinguish between two functions ZFC plays: a largely technical function and a philosophically more substantial function.⁵¹ The technical function of ZFC includes a precise definition of the system's language, a definition of model and truth in a model, a definition of proof, and a precise rendering of our general characterization of formality. All but the last are familiar from standard textbooks in mathematical logic, so I will limit myself here to the precise rendering of formality. One way to reformulate our characterization of formality (concisely) using the precise vocabulary of ZFC is:

(FORMALITY**) *An operator/property is formal iff it is invariant under (all) isomorphisms.*

Due to its precision, we may call this characterization a "criterion" of formality.

Using largely familiar vocabulary, we can explain this criterion in two steps as follows:

⁵¹ It may happen that functions first viewed as purely technical are later, or in a different context, viewed as having philosophical significance. Still, at any particular time and context it might be useful to introduce this division.

Step 1: Explanation of isomorphism. Two structures of objects are isomorphic iff there is a 1-1 and onto function f from the universe of one to the universe of the other⁵² such that f is “structure preserving”, i.e., the second structure is the “image” of the first structure under f . To see what this means let us go back to an example we gave in discussing the non-specialized notion of invariance under 1-1 replacements of individuals. Let A be a set of ten distinct men and ten distinct women: $A = \{m_1, \dots, m_{10}, w_1, \dots, w_{10}\}$ ⁵³, and consider the structure $S1 = \langle A, \beta \rangle$ where β is a subset of A , say the set of all the men in A . Now, consider another structure, $S2 = \langle A', \beta' \rangle$ where A' is a set of ten distinct even numbers and ten distinct odd numbers— $A' = \{e_1, \dots, e_{10}, o_1, \dots, o_{10}\}$ —and β' is the set of all the even numbers in A' . $S1$ and $S2$ are isomorphic iff there is a 1-1 and onto function from A to A' , such that $S2$ is the image of $S1$ under this function (i.e., A' is the image of A and β' is the image of β under this function). It is easy to see that there is such a function; therefore, $S1$ and $S2$ are isomorphic. Now consider the structures: $S3 = \langle A'', \beta' \rangle$, $S4 = \langle A', \beta'' \rangle$, and $S5 = \langle A', \beta''' \rangle$, where A' and β' are as before, $A'' = \{e_1, \dots, e_{10}, o_1, \dots, o_9\}$, $\beta'' = \{e_1\}$, and $\beta''' = \{e_1, \dots, e_9\}$. We readily see that in all these cases there is no function satisfying the conditions stated above. Therefore, $S1$ is not isomorphic to $S3$, $S4$, or $S5$.

In the general case a structure S is an $(n+1)$ -tuple, $\langle A, \beta_1, \dots, \beta_n \rangle$, where A is a non-empty set of individuals (the universe of the structure) and each of the β_1, \dots, β_n is either a member of A or a construct of members of A (an n -tuple of members of A , a subset of A , an n -place relation on A , etc.). Structures S and S' are isomorphic iff there is a 1-1 and onto function from A to the universe of S' such that S' is the image of S under this function.

Step 2. Explanation of formality as invariance under isomorphisms. To determine whether an n -place operator O (property P) is formal, we consider structures S of the form $\langle A, \beta_1, \dots, \beta_n \rangle$ where β_1, \dots, β_n , in that order, are arguments of $O(P)$. $O(P)$ is formal iff for any isomorphic structures $S1 = \langle A, \beta_1, \dots, \beta_n \rangle$ and $S2 = \langle A', \beta'_1, \dots, \beta'_n \rangle$ of this kind, O assigns the same truth value to $\langle \beta_1, \dots, \beta_n \rangle$ in A as to $\langle \beta'_1, \dots, \beta'_n \rangle$ in A' ($\langle \beta_1, \dots, \beta_n \rangle$ has the property P in A iff $\langle \beta'_1, \dots, \beta'_n \rangle$ has the property P in A').⁵⁴

Thus, consider the first-level properties *is-a-man* and *is-self-identical*, and the second-level properties *is-a-non-empty-property* (\exists) and *is-a-property-of-humans*. It is easy to see that *is-self-identical* and *is-a-non-empty-property* (\exists) are formal according to our precise characterization, while *is-a-man* and *is-a-property-of-humans* are not, just as they were according to our non-specialized characterization. Using a

⁵² f from A to A' is 1-1 iff f does not assign to any two objects in A the same object in A' . f from A to A' is onto iff every object in A' is assigned by f to some object in A .

⁵³ In the earlier example (p. 227) we used “D1” for what we now call A .

⁵⁴ (i) To prevent misunderstanding, let me note that the meaning of “ P ” here is fixed, and the question is whether $\langle \beta_1, \dots, \beta_n \rangle$ has P (in A) and whether $\langle \beta'_1, \dots, \beta'_n \rangle$ has P (in A'). I will explain why we relativize these questions to A/A' shortly.

(ii) Notation: In examples I will sometimes use the notation “ $\langle A, \beta \rangle$ ” for “ $\langle A, \beta \rangle$ ”.

specialized characterization makes these results precise as well as extendable to less familiar operators/properties.

Now, one question that might be asked is why we need to take a domain (universe) into account in determining whether a property (operator) is formal. The answer is that we need to do so because in the case of some properties this is relevant. Compare the second-level properties *is-a-non-empty-property* (\exists) and *is-a-property-of-all-but-ten-things*. The former is not relative to an underlying universe but the latter is. Similarly, properties such as *is-a-universal-property* (\forall) and *is-a-property-of-most-things* are. In the interest of unity we formulate the criterion as a whole in terms that involve a universe. This does not affect the treatment of operators/properties which are not sensitive to domains, like non-emptiness (\exists) or identity ($=$), while allowing us to apply it to a broader array of operators/properties.

The above formulation of the specialized formality criterion is directed at predicative operators/properties. This includes sentential operators/properties (like $\&$) when they behave as predicative operators (for example, in a context like “ $Bx \& Cx$ ”, where $\&$ behaves like \cap). For sentential operators behaving as such—e.g., $\&$ in a context like “ $S_1 \& S_2$ ”, where S_1 and S_2 are sentences whose internal structure is not taken into account—our specialized criterion of formality is the standard Boolean criterion of logicity, familiar from standard textbooks. This criterion accurately captures the idea expressed by our general characterization of formal sentential operators and presents it in mathematical terms.

Turning to *logicality* and *logical constants*, here there is no need for new, specialized formulations. On the specialized reading, we simply understand all references to formality as references to the specialized characterization (criterion) of formality.

The historical roots of our specialized criterion of formality go back to mathematical criteria of logicity proposed by Mostowski (1957), Lindström (1966), and Tarski (1966).⁵⁵

So far we have described the largely technical role ZFC plays in our specialized account of logic. Let us now turn to its more substantial role, namely that of answering specific yet philosophically significant questions that our general account leaves open. One of these questions concerns the totality of models. We said that a logical system must be so constructed that the totality of its models represents the totality of formally-possible situations vis-a-vis its language (or any of its given languages). But we left it open what the precise extent of this totality is. We also left open the related question of the precise totality of isomorphic structures under

⁵⁵ It is hard to attribute a philosophical motivation to these papers since they say virtually nothing of philosophical import about these criteria. (Tarski’s 1966 lecture is not an exception. See Section 10.6, fn. 70.) Two philosophers who explored the philosophical significance of invariance-under-isomorphism as a criterion of logicity before me, yet reached different conclusions, are Peacocke (1976) and McCarthy (1981). I discussed McCarthy in Sher (1991, Chapter 3) and will return to his concerns in Section 7 of this chapter. I will explain my approach to the kind of concerns that Peacocke spoke about—namely, apriority—in Section 10.5.

which formal/logical operators/constants must be invariant. And we left open the question of what the formal laws are, a question that determines what regularities hold in the totality of models and what laws govern the formal operators denoted by the logical constants of a given logical system. These are the kinds of questions to which the choice of a background theory of formal structure matters most. ZFC gives specific answers to these questions, but our general foundational account of logic is not dependent on these answers.

III. A joint perspective

We are now ready to answer the following questions:

1. *Does the current paradigm of a logical system, standard first-order mathematical logic, exemplify the principles delineated in our general account?* The answer to this question is positive. Standard first-order logic can be obtained from our general account by (i) using ZFC as a background theory of formal structure, and (ii) selecting a particular set of admissible logical constants (i.e., constants satisfying our criterion) and building these constants into the system as logical constants in accordance with the guidelines and constraints delineated above. Our formality thesis, thus, together with other elements of our account, offer a theoretical philosophical foundation for standard logic.

2. *Does the current paradigm of a logical system account for all logically valid inferences based on our account?* The answer to this question is negative. For example, the quantifier “Most” (as in “(Most x) Φx ”) is an admissible logical constant on our account. This quantifier denotes the logical/formal operator *Most*, defined as, say, “more than half”⁵⁶. So defined, *Most* is a formal operator (according to our criterion), and as such it is governed by formal laws. As a result, consequences based on this operator (alone or in conjunction with other logical operators) are genuinely logical, and in particular, satisfy the requirement of considerable generality and modal force. But standard first-order logic does not recognize—and indeed cannot detect—these consequences. Therefore, standard first-order logic is “incomplete” as far as its ability to identify all logically valid inferences (or provide resources for drawing all such inferences) is concerned. It takes a family of logical systems to complete the task that standard first-order logic partially performs. Looking at this situation from the point of view of the scope of our foundational account, we may say that although it includes standard first-order logic, it is not limited to it. It extends to so-called “generalized” first-order logics (first-order logic with additional logical constants satisfying our criterion) as well as to second-order logic and other more powerful logics.

⁵⁶ That is, given a universe A and a subset B of A (extension-in- A of some property), *Most* assigns the value True to B iff the cardinality of B is larger than the cardinality of $A-B$.

We have opened this section with the question: What specific features of the world is logic grounded in? Our answer was that logic is grounded in formal laws governing objects and properties in the world. This is our *formality thesis*, which we theoretically developed in this section. The formality thesis has considerable explanatory power. In the next section we will show how it provides an explanation—not just an intuitive explanation, but also a theoretical explanation—of the two central traits we have associated with logic: considerable generality and strong modal force. We will then extend this explanation to other traits attributed to logic in the philosophical literature: topic neutrality, abstractness, basicness, strong normativity, certainty, and (quasi-) apriority.⁵⁷

10.5 From Formality to Generality, Necessity, Topic Neutrality, Strong Normativity, Quasi-Apriority, and More

A. *Formality*. The key to explaining the characteristic traits of logic based on its formality is the distinctive theoretical feature of formality: *invariance under 1-1 replacements of individuals*, or, using specialized terms, *invariance under isomorphisms*. The appeal to invariance should not be surprising since, as noted in Section 10.4, invariance is known for its explanatory power (see, e.g., Woodward 2000 and Nozick 2001). In the case of logic, we explain its traits in terms of the special invariance characteristic of formality: invariance under 1-1 replacements of individuals (both within and across domains).

To see how this theoretical feature explains the characteristic traits of logic, let us go back to the observation that properties in general are *selective* in nature, i.e., “pay attention” or “are attuned” to *some* features of objects but not to others. For example: the property *is-a-person* is not attuned to differences in gender; physical properties (as distinct from, say, biological properties) are not attuned to differences between animate and inanimate objects; etc. This introduces the possibility of partially ordering properties according to the differences between objects they are not attuned to and of using *invariance* as a basis for such an ordering: *is-a-person* is attuned to fewer differences between objects than *is-a-woman* and physical properties are normally attuned to fewer differences between objects than biological properties. This is reflected in the fact that *is-a-person* is *invariant under* more 1-1 replacements of individuals than *is-a-woman* and that physical properties are invariant under more 1-1 replacements of individuals than biological properties. We thus say that the *degree* of invariance of *is-a-person* is *higher than* that of *is-a-woman* and that

⁵⁷ I have not included analyticity in this list. This is because analyticity, as traditionally understood, is not a feature of logic on our account, since it denies logic a significant grounding in reality. Of course, if analyticity is understood as I understand formality, then my account explains the analyticity of logic as well.

generally, the degree of invariance of physical properties is greater than that of biological properties. Generalizing, we may say that the degree of invariance of physics as a field of knowledge is greater than that of biology as a field of knowledge. For the present purpose, however, there is no need to determine a full partial-ordering of properties or fields of knowledge in terms of invariance. All we need is to compare the degree of invariance of formal properties to the degree of invariance of other properties, and, based on this, compare the degree-of-invariance of logic to that of other disciplines.

It is readily seen that the degree of invariance of logic is greater than that of most other disciplines: logical properties (operators, constants), being formal, are invariant under *all* 1-1 replacements of individuals, whereas most other properties are invariant only under *some* 1-1 replacements. For example, neither physical properties nor numerical properties (properties of numbers as individuals) are invariant under 1-1 replacements of stones by numbers, but formal properties are: the property *is-self-identical* does not distinguish between a stone and a number, and likewise, the property *has-cardinality-THREE* does not distinguish between a set, or a property, of stones and its image under any 1-1 replacement of stones by other individuals. We may characterize “greater degree of invariance” by:

The larger the range of individuals among which X—a given property, a field of knowledge, a law, a theory—does not distinguish, the greater its degree of invariance.

Now, the relation *greater-than* can be measured in (at least) two ways: based on *cardinality* and based on *inclusion*. (The set {1,2,3} is greater, or larger, than both {2,3} and {4,5} in terms of cardinality, but it is not larger than {4,5} in terms of inclusion.) The relation of *has-a-greater-degree-of-invariance* that we will use here is based on *inclusion*. And it is readily seen that the above characterization is preserved under this conception of *greater-than*. So are the conclusions that (i) formal properties have a greater degree of invariance than most other properties and (ii) logic has a greater degree of invariance than most other disciplines.

A simple but powerful principle connecting degree of invariance to the traits of logic is:

If X does not distinguish between Y and Z, it applies equally to Y and Z (i.e., if it applies to Y, it also applies to Z).

This principle is especially useful when dealing with laws. If logical laws do not distinguish between physical and mathematical domains, they are equally applicable to both.⁵⁸

⁵⁸ (i) The above principle may be viewed as implying that logical laws apply either to all fields of knowledge or to none. Since we have already discussed the reality of formal properties and laws and given examples of the applicability of formal laws to the physical domain (Chapter 8, Section 4 and, indirectly, earlier in the present chapter), there is no need to consider the second option here. However, there are some subtle points concerning the first option. These will be discussed in Sections 10.9 and 10.10.

B. *Generality and Modal Force.* The strong degree of invariance of formal, hence logical, laws has two dimensions—a *factual* dimension and a *counterfactual* dimension. Its factual dimension concerns actual objects and situations, its counterfactual dimension—counterfactual objects and situations. The former explains the *generality* of formal laws: why they apply to objects and situations in *all fields* of knowledge. The latter explains their considerable *modal force*: why they apply to *counterfactual* objects and situations. The point is that since formal laws have a greater degree of invariance than most other laws, their *factual domain* properly includes the factual domains of most other fields of knowledge, and their *counterfactual domain* properly includes that of most other laws. That is to say, logic, as a theory of inference based on formal laws, is more general and has a greater modal force than most other fields of knowledge.

Let me explain in greater detail how the strong invariance of logic is related to its strong modal force.⁵⁹ The strong invariance of logic consists in the fact that the logical constants and the formal properties and relations they denote do not distinguish between any isomorphic structures of objects, no matter what their elements are. In particular, they do not distinguish between any *actual* and *counterfactual* structures which are isomorphic to each other.⁶⁰ Take the first-order existential quantifier. This logical quantifier denotes the second-level property of non-emptiness, and this property *behaves in exactly the same way* in all isomorphic structures, regardless of whether they are actual or counterfactual, hence in all the situations they represent, regardless of their actuality. The satisfaction conditions of the existential quantifier do not vary from the actual situation, in which the speed of light is 186,000 miles per second, to any isomorphic counterfactual situation in which the speed of light is different from 186,000 miles per second; they do not vary from the actual world, in which there are physical objects, to any isomorphic counterfactual world in which there are only abstract objects, and so on. Because for each actual situation there are isomorphic counterfactual situations, we can say

- (ii) Here I am talking specifically about domains of fields of knowledge, where knowledge is understood as knowledge of the world (or of some aspect of the world). This includes domains of posits which are systematically representative of real objects, like domains of numerical individuals on our account (see Chapter 8, Section 4), but not domains of objects in artistic fiction. I will briefly discuss the case of artistic fiction in Section 10.10.

⁵⁹ For the sake of simplicity I will use standard set-theoretic terminology here.

⁶⁰ By an “actual structure” I mean a structure with *actual* individuals and their actual properties (relations, etc.), and by a “counterfactual structure” I mean a structure which either has some non-actual individuals or has some actual individuals with properties that they do not have in the actual world. An example of an actual structure is the structure $S = \langle A, B \rangle$, where A is a non-empty set of actual photons and B is the set (property-in-extension) of all photons in A whose speed is 186,000 miles per second. An example of a counterfactual structure is the structure $S' = \langle A', B' \rangle$, where A' is, again, a non-empty set of actual photons and B' is a non-empty set of all photons in A' whose speed is 200,000 miles per second. S' represents a counterfactual situation in which the speed of at least some photons is different from what it actually is. Another example of a counterfactual structure is the structure $S^* = \langle A^*, B^* \rangle$ where A^* is a non-empty set that contains Martians, and B^* is the property of, say, being green.

more generally that the existential quantifier behaves in the same way regardless of whether a given situation is actual or counterfactual. And this means that the laws governing the non-emptiness property have an extremely strong modal force, spanning an enormous array of counterfactual situations. (Indeed, since the non-emptiness property is defined over any universe of any cardinality, whether this cardinality is instantiated in the actual world or not, these laws govern all counterfactual situations, so long as they have a universe of a determinate cardinality.) This, in turn, means that all logical consequences grounded in the formal laws governing the existential-quantifier property have an extremely strong modal force.

All this generalizes to consequences grounded in any formal laws. Since all valid logical consequences are grounded in such laws, all valid logical consequences have an extremely strong modal force, and the source of their modal force is the extremely strong invariance of the formal properties (denoted by the logical constants) which these laws govern.⁶¹

C. *Topic Neutrality*. We have already explained why the topic neutrality of logic does not mean that logic does not have a topic of its own. Topic neutrality, in the sense applicable to logic, has to do with scope or range of applicability: logic is topic neutral iff it applies to all fields of knowledge equally, regardless of what *their* specific subject matter is. That logic satisfies this condition follows from its generality, which, as we have just seen, follows from its formality. More directly, since logical operators, being formal, do not distinguish between arguments belonging to different fields of knowledge, they apply to all fields, regardless of *their* “topic”.

D. *Abstractness*. To abstract is to overlook, be indifferent to, not to distinguish, features and regularities of objects that are distinguished prior to abstraction. To be abstract, in the context of knowledge, is to abstract from causal, spatiotemporal, observable features and regularities—to abstract from the “concrete”. Formal properties abstract from features of objects whose degree of invariance is weaker than theirs, and concrete features fall under this category. As such, formal features, their laws, and the logical laws grounded in them, are highly abstract. Although formality is not necessary for abstractness, it is sufficient for it. In particular, logic’s formality ensures, and explains, its abstractness.

E. *Basicness*. Logic is traditionally thought to be a “basic” discipline, and this observation retains its validity in our holistic approach. Of course, there are different ways of being basic. Logic is basic in the sense that *all (or most) disciplines are bound by its laws* whereas logic itself is not bound by their laws. This, too, can be explained by the formality of logic. Due to its formality—strong degree of invariance—logic’s

⁶¹ I should indicate, however, that while the formality of logic implies (not logically, but in a broader philosophical sense) that logic is modally stronger and more general than most other disciplines, it does not imply that logic is general or necessary in an *absolute* sense. See discussion in Section 10.7. Other issues concerning the modal force of logic are also discussed in that section.

scope of applications includes all other disciplines, but not vice versa. As a result, all other disciplines must obey the laws of logic, but logic is (for the most part) not obliged to obey theirs.⁶²

Another way in which logic is basic has to do with its ubiquity. This point we already made in connection with the *logicality principle of truth* (Chapter 8, Section 5). There is an important sense in which formal notions, and the logical notions based on them, are pervasive in human thought and discourse. There is a sense in which we cannot make any significant move in thought and knowledge without making extensive use of formal notions. We cannot think significantly about objects without thinking of them as *different* from other objects, as being in the *complement*, *union*, *intersection* of some properties and relations, and so on. We cannot think about properties without thinking of them as being *empty* or *universal*, as *holding of n or m objects*. We cannot think about relations, especially scientific relations, without thinking of them as exemplifying formal patterns (reflexivity, symmetry, transitivity, ...). And so forth. Looking at the world, we see formal features everywhere, and talking about the world, we constantly use expressions that refer (directly or indirectly) to these features. To the extent that regimenting the world in formal terms is so basic to humans, logic, as a theory of formal reasoning, is equally basic.⁶³

F. *Strong Normativity*.⁶⁴ The strength of logic's normativity can also be explained by its formality. In fact, it is the mirror image of its basicness.⁶⁵ Due to its especially strong invariance compared with other disciplines, these disciplines are bound by the laws of logic, but logic is not bound by theirs. Since chemical properties, for example, are not preserved under isomorphisms, logic does not distinguish chemical differences between objects and, as a result, is not subject to laws governing chemical properties. But chemistry does distinguish formal differences between objects. For example, it distinguishes between a molecule that has *one* atom of a certain kind and a molecule that has *two* atoms of that kind. It distinguishes between chemicals that have *both* feature A *and* feature B and those having feature A *but not* feature B. And so on. Since the scope of formal laws properly includes that of chemical laws, chemical theories have to take into account, and indeed obey, formal, hence logical laws. But not the other way around.

Indeed, the normative strength of logic, like its basicness, has two aspects. To the extent that all disciplines are governed by logical norms, logic has an especially strong

⁶² Adding the qualifier "for the most part" is due partly to the holistic nature of our account and partly to the point noted in the last footnote, namely that logic is highly, but not "absolutely", general or necessary, a point that will be explained in Section 10.7.

⁶³ See also Chapter 5, Section 2 (abstract features of reality) and Chapter 8, Section 4 (mathematics).

⁶⁴ Here we are concerned with the *strength* of logic's normativity. In the next section we will discuss the *source* of its normativity.

⁶⁵ We can regard logic as relatively highly normative without enforcing a rigid, one-dimensional ranking of disciplines with respect to their normativity.

normative force; and to the extent that formal properties, and the logical expressions denoting them, are pervasive in all areas of knowledge, logical norms are constantly exercised. That is to say, logic is highly normative both in the sense that its norms have a high degree of normativity and in the sense that the volume of use of its norms is very high. Both of these are (or at least can be) explained by its formality.

G. *Certainty*. Logic is thought to be highly certain. It is of course not highly certain in the sense that one is unlikely to make errors in applying its laws—people make errors in logical reasoning quite often. But logic is *relatively certain* in the sense that its laws are less likely to be refuted by new discoveries, and especially by empirical discoveries, than laws of other disciplines. This does not mean that logic is immune to new discoveries altogether, either in its own area (recall Russell's discovery of a paradox in Frege's logic), or in related areas (due to the interconnectedness of fields of knowledge). But it means that in a certain deep sense logic is more stable than other disciplines and *its laws are more shielded from new results*—especially empirical results—than other laws. Once again, this can be explained by its formality. Due to their strong invariance, formal, hence logical, operators are indifferent to most aspects of reality (to most features of, and variations between, objects), and as a result, the laws governing these operators are not affected by investigations concerning those aspects. Empirical discoveries might undermine our judgments about how many objects possess a certain physical property, or which, and how many objects stand in the intersection of certain physical properties. But they are unlikely to affect the formal judgments that are relevant to logic—e.g., judgments about the formal laws of cardinality or the formal laws of intersection. This leaves logic relatively immune to empirical threats, and in this sense logic is highly (or relatively highly) certain.⁶⁶

H. *Quasi-Apriority*. Traditionally, logic is characterized as purely apriori; on our view, it is *quasi-apriori*, in the sense explained in Chapter 5, Section 3. Either way, the apriority of logic can be explained based on its formality. If you view our knowledge of the formal as apriori, you will be able to use the formality of logic to explain its apriority; if you view this knowledge as quasi-apriori, you will use the formality of logic to explain its quasi-apriority. The explanation itself can be either direct or indirect. A direct explanation will proceed along similar lines

⁶⁶ (i) As in the case of its other traits, it is important to note that the claim is that logic is *relatively* immune to empirical discoveries (hence *relatively* certain), not that it is *completely* immune (or *absolutely* certain). Some ways in which logic, whose knowledge is largely intellectual, might be affected by empirical discoveries will be discussed in Section 10.9.

(ii) It is important to distinguish the claim that logic is certain from the claim that logic is safe from controversy. Non-controversiality is by no means a feature a logical theory has to have in order to fulfill its designated role in knowledge. Whether a theory is controversial or not is a matter of how people in the field react to it, and as the history of science, mathematics, and philosophy teaches us, there is no direct correlation between being controversial and being false, unjustified, or uncertain (in our sense).

to those used to explain the certainty of logic: Due to its strong invariance, the formal is indifferent to most aspects of reality, and this includes those aspects that can be studied/established using sensory perception or, more generally, the *empirical method*. Therefore, our theory of the formal—and, by extension, logic—is largely/completely *non-empirical*, i.e., *intellect-based*.⁶⁷ An indirect explanation will be based on the necessity, abstractness, and generality of logic: logic is too necessary, too abstract, and too general to be studied, or be established, by empirical methods.⁶⁸

10.6 Source of the Normativity of Logic, Tarski's Problem, Truth and Logical Truth, and Other Issues

We have seen how the formality thesis explains the grounding of logic, or more specifically, what aspect of reality logic is grounded in. We have further seen how it explains logic's ability to provide a highly general and modally strong method of inference. And we have also seen how it explains the other characteristic traits of logic, from its topic neutrality to its quasi-apriority. Further contributions of the formality thesis to the explanatory power, problem solving capacities, and fruitfulness of the present account pertain to (A) an understanding of the source of the normativity of logic, (B) a solution to Tarski's 1936 problem, (C) a systematization of logical constants, (D) a critique of the "standard first-order thesis", (E) developments in mathematics and linguistics, and (F) an understanding of the relation between truth and logical truth. Let us now turn to these contributions.

A. The Source of the Normativity of Logic. We have explained why the normativity of logic is stronger than that of other disciplines. But what is the *source* of logic's normativity? In approaching the normativity of logic philosophers have taken two approaches: an *agent-oriented approach* and a *theory-oriented approach*. These two approaches give rise to different questions, different explanations, different constraints, etc. The agent-oriented approach is characteristic of such works as

⁶⁷ The apriorist will say that it is *completely* non-empirical, the quasi-apriorist—that it is *largely* non-empirical. What the latter means will be explained in Section 10.9. See also discussion of the quasi-apriori in Chapter 5, Section 3 and discussion of holism in Chapter 2, Section 2 and Chapters 3 through 6.

⁶⁸ (i) Pure apriorists can give an explanation very similar to ours. Due to its strong invariance, the formal is indifferent to *all* aspects of reality that can be studied using sensory perception or, more generally, the empirical method. Therefore, our theory of the formal—and, by extension, logic—is *apriori*.

(ii) Although formality can be used to explain both the apriority and the quasi-apriority of logic, it is a significant aspect of our overall theory that logic is quasi-apriori rather than either purely apriori or empirical. By viewing logic as quasi-apriori we avoid a conflict between our rejection of the (pure) apriority of logic and our rejection of the view that logic is empirical. This is something I did not sufficiently emphasize in my earlier writings, causing Hanson (1997) to misunderstand my view. Hanson, however, later recognized that there need not be a conflict between these two views (see Hanson 2002: 250–1).

Harman (1986, 2009) and Field (2009a). The theory-oriented approach is the one I take here. The agent-oriented approach focuses on how actual agents in the world engage in logical reasoning, whether logical principles are sufficient to justify revision of a given belief by an agent, how various biological and psychological problems affect logic's normativity for agents, etc. The theory-oriented approach views normativity from a different angle and asks questions that are usually left aside in the agent-oriented approach: What is the source of the normativity of theories in general and of logic in particular? Does the normativity of the logical method depend only on the mind or also on the world? How are disciplines in general, and logic in particular, normative for other disciplines? And so on. Here we will focus on the first question. Our answer to this question is divided into two parts: (a) The source (or an important source) of the normativity of any field of knowledge is *truth*; (b) The source (an important source) of the normativity of logic is *formal truth*.

(a) *Truth as a source of epistemic normativity in general.* There is an important sense in which truth is a *norm*. This sense was expressed in our view (in Chapter 8, Section 1) that truth is primarily a *standard* for sentences and theories and only secondarily a property. Following Williams (2002), we may call this norm "truthfulness". Here we are primarily interested in truthfulness as a norm of assertion and theorizing. Since the truth of our theories is primarily a matter of their agreement with reality, we may say that, indirectly, reality itself, and in particular, its laws, exert normative force over our cognitive activities. In Frege's words:

Any law asserting what is, can be conceived as prescribing that one ought to think in conformity with it, and is thus in that sense a law of thought. This holds for laws of geometry and physics no less than for laws of logic (Frege 1893: 12).

The norm of speaking truthfully lies in the intersection of ethics and epistemology. It says that we have a moral obligation to aim at the truth in epistemic contexts. MacFarlane's explanation of Frege's view further clarifies the objectual underpinnings of this normativity:

On Frege's view, ... it is a feature of *all* descriptive laws [that they are normative.] ... [C]onsider the "game" of thinking about the physical world (not just grasping thoughts, but evaluating them and deciding which to endorse). ... "[M]oves" in this game—judgments—can be assessed as correct or incorrect. Judgments about the physical world are correct to the extent that their contents match the physical facts. Thus, although the laws of physics are descriptive laws—they tell us about (some of) these physical facts—they have **prescriptive consequences for anyone engaged in the "game" of thinking about the physical world**: such a thinker *ought not* make judgments that are incompatible with them. Indeed, insofar as one's activity is to count as making judgments about the physical world at all, it must be assessable for correctness in light of the laws of physics. In this sense, the laws of physics provide *constitutive norms* for the activity of thinking about the physical world (MacFarlane 2002: 36–7).

This suggests that different laws might have different scopes of normativity. Since the laws of logic are universal, logic is universally normative. In Frege's words:

From the laws of [logic] there follow prescriptions about asserting, thinking, judging, inferring [in general] (Frege 1918: 1).

One result of these considerations is a close connection between the descriptive and the normative:

[A]lthough logical laws are [descriptive and] not prescriptive in their content, they *imply* prescriptions... For example, [they imply that] one ought not believe both a proposition and its negation. Logical laws, then, have a dual aspect: they are descriptive in their content but imply norms for thinking (MacFarlane 2002: 36).

(b) *Formal truth as the source of logical normativity.* Although the general source of logic's normativity is the same as that of other disciplines, namely, truth, its normativity is grounded in a special type of truth. On our account, it is grounded in *formal* truth. Take, for example, the logical norm: "Do not believe both a proposition and its negation". This norm is grounded in the formal law that no object lies in the intersection of any property and its complement in a given universe (or its analog for states of affairs and being/not-being the case).

So far we have seen three features of logical normativity that distinguish it from the normativity of disciplines like physics: first, logical normativity is formal; second, logical normativity is universal; and third, logic is not subject to the norms of physics while physics is subject to the norms of logic. There is, however, an additional difference between the normativity of logic and that of physics. Unlike physics, logic carries its normativity on its sleeve, so to speak. This is due to the fact that logic does not talk directly about reality: it talks about linguistic entities used in talking about reality. In this respect the normativity of logic is similar to that of grammar. But in being veridical, i.e., in being concerned with truth and being grounded in reality, logic resembles physics more than it resembles grammar. The normativity of logic, like that of physics and unlike that of grammar, is grounded primarily in truth rather than in convention.

B. Solution to Tarski's 1936 Problem. Both historically and substantively, one of the cornerstones of contemporary logic is Tarski's 1936 definition of logical consequence. But this definition, as Tarski pointed out, is critically dependent on the existence of a criterion for logical constants:

Underlying our whole construction is the division of all terms of the language discussed into logical and extra-logical. This division is certainly not quite arbitrary. If, for example, we were to include among the extra-logical signs the implication sign, or the universal quantifier, then our definition of the concept of consequence would lead to results which obviously contradict ordinary usage (Tarski 1936a: 418).

This is a deep philosophical problem, but in 1936 Tarski did not see his way to a solution to this problem:

[N]o objective grounds are known to me which permit us to draw a sharp boundary between the two groups of terms (Tarski: 418–19).

He concluded his paper on a skeptical note:

Further research will doubtless greatly clarify the problem which interests us. Perhaps it will be possible to find important objective arguments which will enable us to justify the traditional boundary between logical and extra-logical expressions. But I also consider it to be quite possible that investigations will bring no positive results in this direction (Tarski: 420).⁶⁹

Our account solves this problem by identifying logical constants with constants denoting formal operators and using invariance-under-isomorphisms as a criterion of formality. Significantly, the considerations that led us to this solution directly address Tarski's main concern in his (1936a) paper, namely, yielding an adequate notion of *logical consequence*. Two conditions that an adequate notion of logical consequence has to satisfy, according to Tarski, are *necessity* and *formality*. We may express these conditions as follows: An adequate notion (or definition) of logical consequence is one according to which:

- (a) X follows *logically* from K *only if* X follows *necessarily* from K, and
- (b) X follows *logically* from K *only if* X follows *formally* from K.

The Tarskian challenge, therefore, is to find a criterion for logical constants that, when used in conjunction with the model-theoretic definition of logical consequence, yields consequences that satisfy the philosophical conditions of necessity and formality.

Our formality thesis directly meets Tarski's challenge. Invariance under isomorphisms is an objectual criterion of formality, and this kind of formality, as we have seen in the last section, ensures necessity. Furthermore, the objectual formality of logical operators serves as a basis for their syntactic formality—for their being chosen as markers of logical form—so that even if what Tarski had in mind is syntactic formality, our solution provides a philosophical underpinning for his view. Rather than viewing logical form as based on an arbitrary choice of logical constants, we

⁶⁹ Readers familiar with Etchemendy (1990) will remember his claim that Tarski's 1936 problem of logical constants is a "red herring". However, Etchemendy regards the problem of logical constants as a red herring not because he thinks logical constants do not pose a genuine problem to Tarski's definition, but because he thinks that Tarski's definition is plagued by other problems as well, so merely solving the logical-constants problem is not sufficient to establish its adequacy. Accordingly, the question of a criterion for logical constants is a red herring only if the other problems Etchemendy alludes to are real. In my view, they are not. As I explained in Sher (1996), the view that they are real is based on a mistaken construal of Tarskian semantics, namely on its construal as an "interpretational" semantics, and more basically, on a mistaken "either-or" assumption, namely on the assumption that Tarski's semantics is either "interpretational" or "representational". On the present construal of Tarskian semantics it is neither interpretational nor representational (in Etchemendy's sense), and therefore none of the problems raised by Etchemendy arise for it. As a result, the significance of Tarski's logical-constants problem is not mitigated.

view it as based on a non-arbitrary choice of constants, namely, one that selects logical constants according to their denotation, where permissible denotations are ones that ensure the *necessity* and *formality* of the ensuing logical consequences.⁷⁰

C. Contribution to the Systematicity of Logic. The formality criterion closes a gap between the treatment of logical operators in sentential and predicate logic. On the one hand, the standard criterion for logical operators in sentential logic is a theoretically systematic and informative criterion, namely, the Boolean or truth-functional criterion. On the other hand, the standard criterion for logical operators in predicate logic is theoretically uninformative and unsystematic, namely, a criterion by enumeration: C is a logical operator iff C is a logical operator of sentential logic, or C is " \forall " / " \exists ", or C is " $=$ ", or C is definable from any of the operators in the above list. The invariance-under-isomorphisms criterion does for predicate logic what the standard Boolean criterion does for sentential logic. It identifies a distinctive trait of logical operators on the predicative level just as the traditional criterion does on the sentential level. Furthermore, both of the standard criteria are sometimes viewed as merely technical, but the formality thesis invests them with philosophical content and significance.

D. The "Standard First-Order Thesis": A Critical Perspective. For a long time, the prevalent view on logic was that logic is standard first-order logic; that is, the fundamental principles of logic are fully captured by the standard system. Barwise (1985) calls this view "the first-order thesis". A more appropriate name would be "the *standard* first-order thesis", which is the view that first-order logic is limited to the "standard" logical constants—truth-functional connectives, identity, the universal/existential quantifier, and the constants defined in terms of these. In

⁷⁰ A historical note: in his 1966 lecture Tarski offered a criterion for logical operators that can be viewed as his own solution to his 1936 problem. However:

- (i) Tarski did not present his criterion as a solution to this problem, and indeed did not mention the 1936 problem (or paper) at all in his 1966 lecture. Furthermore, Tarski explicitly discouraged us from attributing any philosophical significance to his criterion concerning the question "What is logic?". So one has to be cautious about attributing to Tarski in 1966 intentions he may not have had at that time.
- (ii) Tarski's criterion for logical operators is *invariance under permutations*. This criterion is superficially similar to our own criterion (*invariance under isomorphisms*) but, as noted by McGee (1996), it is not a suitable candidate for a criterion of logicity. Consider an operator that behaves like the existential quantifier in universes of apples and like the universal quantifier in universes of oranges. Such an operator is logical according to Tarski's criterion but not according to ours. Clearly, such an operator should not be counted as logical, distinguishing, as it does, features of objects that are not relevant to logic (being an apple, being an orange). It should be noted, however, that due to differences between Tarski's background theory (Russellian type-theoretic logic) and the current background (contemporary set theory, e.g., ZFC) this deficiency might disappear.
- (iii) If we do look for a philosophical idea underlying Tarski's criterion, then, as noted by Bonnay (2008), it is the idea that logical consequence is the most *general* consequence there is. However, viewing Tarski's criterion (or ours) as a criterion of generality is misguided, as we will see in Section 10.7.

spite of the fact that the standard first-order thesis is still widely accepted by philosophers, very few theoretical justifications have been given to it. Commonly, it is simply taken for granted or arrived at by “following the experts”. Sometimes, its acceptance is justified based on pragmatic considerations, as by Quine (1970/86), who says the Gödel-completeness of standard first-order logic is pragmatically desirable.⁷¹ Sometimes, it is justified (especially by mathematicians) by the “nice” metalogical properties of this logic, like the Löwenheim-Skolem property.⁷² But these considerations are less than compelling. It is well known that other logics—e.g., first-order logic with the generalized quantifier Uncountably-Many (Keisler 1970)—are also complete. And “nice” features, like the Löwenheim-Skolem feature, have an “ugly” face, marking, as they do, serious limitations of standard first-order logic, e.g., inability to handle inferences based on infinitistic logical notions (like “uncountably many”). All these suggest that we approach the standard first-order thesis in a critical spirit.

Of course, there have always been philosophers who approached standard first-order logic with a critical spirit. But the advent of a systematic criterion of logical constants (by Lindström and others) led Barwise (1985) to question the standard choice of logical constants. Why, with the discovery of a rich trove of admissible logical constants, Barwise asks, do logicians go on limiting themselves to those chosen long before these discoveries? Our formality thesis adds a new, philosophical dimension to Barwise’s question: Why, if richer first-order logical systems, incorporating new logical constants, produce logical consequences which are just as modally strong, general, and formal as the standard ones, do logicians continue to limit themselves to the standard systems?

E. Fruitfulness in Mathematics and Linguistics. Our specialized invariance criterion, as originally developed by Mostowski, Lindström, and their followers, led to important developments in mathematics and linguistics. The former include the development of new fields of mathematical logic such as model-theoretic logics and abstract logic (see e.g., Barwise and Feferman 1985). Among the fundamental results are Lindström’s (1969, 1974) characterizations of standard first-order logic as the only model-theoretic logic (essentially, the only logic sanctioned by our specialized criterion) that has certain metalogical properties and Keisler’s (1970) and others’ discovery of Gödel-complete non-standard logics. The latter include major developments in linguistic semantics, centered around “generalized” quantifiers, and pioneered by Higginbotham and May (1981), van Benthem (1983), Westerståhl (1985),

⁷¹ Gödel’s completeness theorem says that in standard first-order logic the proof method and the semantic method give the same results: $\Sigma \vdash \sigma$ iff $\Sigma \models \sigma$.

⁷² A logic L has the Löwenheim-Skolem property iff for every sentence (or set of sentences) X of L , if X has a model then X has a countable model. In particular, if X has an infinite model it has a countable model.

Keenan and Stavi (1986), and many others. (For an up-to-date survey see Peters and Westerståhl 2006.)

F. *Truth and Logical Truth.* In developing an account of truth in Chapter 8, we introduced the “logicality principle of truth”. This principle says that *logical structure* is a *universal* determinant of truth (at work in all fields of knowledge), yet only a *partial* determinant of truth (one among many). In introducing this principle we proclaimed that Tarski’s theory of truth makes a major contribution to our understanding of the logicality principle, yet it leaves much philosophical work to be done. This includes (i) an informative account of logical constants—the main constituents of logical structure, (ii) an explanation of how, given that truth is based on correspondence, the logical portion of the truth conditions of sentences is based on correspondence, and (iii) a related explanation of how the truth of logical truths is based on correspondence.

(i) *Informative account of logical constants.* The formality thesis offers a substantive account of logical constants: their denotations, their formality, how their formality enables them to give rise to highly general and modally strong consequence, how they are built into a syntactic-semantic logical system, what their role in proof is, and so on.

(ii) *Explanation of how the contribution of logical structure to truth is based on correspondence.* Briefly: logical constants denote objectual properties of a certain kind (formal properties) and as such have objectual satisfaction conditions. Accordingly, the satisfaction conditions of the logical structure of a given (logically structured) sentence are also objectual, connecting the truth value of that sentence systematically to some formal aspect of the world.

Thus, take any sentence with the logical form:

$$(37) \quad (\forall x)(Bx \vee Cx).$$

Using “B” and “C” for the properties denoted by “B” and “C” respectively (or for the extensions of “B” and “C” in a given universe of discourse), the truth conditions of (37), on our account, can be stated as follows:

“ $(\forall x)(Bx \vee Cx)$ ” is true
iff
every individual in the world (or in a given universe of discourse) stands in the **union**
of B and C in the world (or in that universe of discourse),
or iff
this **union** is **universal** in the world (exhausts the given universe of discourse).

This condition exhibits a systematic connection between what (37) says about the world and the formal features of the situation that has to hold in the world (or in a given domain of objects in the world) for the sentence to be true, and as such it is a bona fide correspondence condition.

(iii) *Explanation of how the truth of logical truths is based on correspondence.* Since the truth of logical truths is based primarily on their logical structure, and since the truth conditions associated with a given logical structure are, as we have just seen, correspondence conditions, the truth of logical truths is based on correspondence. Indeed, our account says more than this on logical correspondence. The truth of logical truths is based on correspondence not just with any formal facts about the world, but with *formal laws* governing the world. This is because our semantic definition of *logical truth* is the Tarskian definition “ σ is logically true iff σ is true in every model”, and formal facts that hold in all formally possible situations—the situations represented by Tarskian models, on our account—are *formal laws*. Thus, consider the logical truth

$$(38) \quad (\forall x)(Bx \vee \sim Bx).$$

Its truth condition, according to our account, is:

“ $(\forall x)(Bx \vee \sim Bx)$ ” is logically true
 iff
 “ $(\forall x)(Bx \vee \sim Bx)$ ” is true in every model
 iff
 in every formally possible structure of objects, every individual stands in
 the union of \underline{B} and its complement (in that structure)
 iff
 it is a formal law governing reality that in every domain of individuals, every
 individual stands in the union of \underline{B} and its complement in that domain,
 or iff
 it is a formal law governing reality that in every domain the union of \underline{B} and its
 complement in that domain is universal (exhausts the domain).

This truth condition exhibits a systematic connection between what (38) says about the world and what holds in the world, namely, between what (38) says and a certain formal law governing the world. Comparing this correspondence condition with the correspondence conditions of sentences expressing physical laws, we may say that like the latter, this correspondence condition connects (38) with an objectual condition on the world, but unlike the latter, it connects (38) with a *formal* correspondence condition rather than with a *physical* correspondence condition.⁷³

⁷³ Among the questions a reader might ask at this point are: (a) Does the world as a whole count as a domain on our account? (b) Given that any two logical truths are logically equivalent, do they correspond to the same formal law? (c) How sharp is the distinction between formal and physical laws? I will address the first question in Section 10.7. As for the second question, my answer is negative: being logically equivalent is being true in the same models (representations of actual and counterfactual situations), but different formal laws can hold in the same situations. Concerning the third question: I have not discussed physical correspondence in detail in this essay. (I think it would be more interesting to discuss it in connection with epistemic freedom.) What I can say until then is that on the one hand, there is a

10.7 Questions and Objections: Logical Constants, Invariance, Generality, and Necessity

I. Logical constants and invariance

Some philosophers/logicians object to our characterization of logical constants on various grounds: some (e.g., Hanson 1997 and Gómez-Torrente 2002) object to the very idea of a theoretical criterion of logicity, preferring a looser, pragmatic approach to logic instead. Some (e.g., Dutilh Novaes 2012 and Sagi 2014) question the centrality of logical constants/operators to a theoretical understanding of logic. Some (e.g., Belnap in conversation) say our specialized version of this criterion—invariance under isomorphisms—undergenerates. And some (e.g., Peacocke 1976, McCarthy 1981, and Feferman 1999, as well as Hanson 1997 and Gómez-Torrente 2002) say it overgenerates.⁷⁴

One of the main difficulties in both assessing and responding to these objections lies in the fact that none of the objections addresses any of the theoretical motivations or advantages of our criterion. Nor do they take into account the new methodological tools and other developments that shape our foundational matrix: holistic foundational methodology, dynamic model of knowledge, quasi-apriority, and so on. This creates a certain degree of incommensurability between our support of the criterion and the existent objections to it.

Our constraints on, and characterization of, logical constants are normative and theoretical. Our goal is to characterize logical constants in accordance with their role in a logical system performing a certain theoretical role in our system of knowledge. Specifically, the goal is to characterize the logical constants in such a way that, given the rest of the syntactic-semantic apparatus, logical consequences transmit truth from premises to conclusion universally and with the requisite modal force. This goal is different from other goals that philosophers associate with definitions of logical constants. In particular, our goal is different from that of many of the critics, whose focus is on natural language and its everyday usage. Thus our goal is different both from the goal of describing how people commonly use the notion “logical constant”

substantial difference in counterfactual force between physical and formal laws. But on the other hand, some physical theories are so thoroughly invested in formal laws that there might not be a clear boundary between their laws and (purely) formal laws.

⁷⁴ As noted earlier, I addressed McCarthy’s objection in Sher (1991, Chapter 3). I responded to Hanson’s objection in Sher (2001), to Gómez-Torrente’s objection in Sher (2003), and to Feferman’s objections in Sher (2008, 2013), and I will return to some of these below. For counter-responses see Hanson (2002) and Gómez-Torrente (2003). Peacocke is concerned with apriority. I explained my position toward the apriority of logic in Section 10.5 and towards apriority in general in Chapter 5, Section 3. Another criticism, due to Bueno and Shalkowski (2013), does not distinguish my approach from other purportedly problematic Tarskian approaches. Thus, (i) Bueno and Shalkowski view the model theoretic approach as either Platonist or nominalist, whereas my approach differs from both, (ii) they identify the model theoretic approach with ZFC whereas I offer a general account of this approach that is independent of ZFC, (iii) they treat circularity in foundationalist terms whereas I view it in holistic terms, etc.

in natural language and from the goal of explaining how people grasp the meaning of logical-constant expressions in natural language (see MacFarlane 2005/9). The latter are worthy goals, but they are not our goals, and therefore criticizing our approach from the point of view of these goals misses the point.

Still, there are some things I can say:⁷⁵

A. *Pragmatist vs. Theoretical Approach to Logic.* The overall epistemic theory developed in this essay says that all knowledge, including philosophical knowledge, requires both pragmatic and theoretical resources, and this applies to foundational pursuits as well. The extreme pragmatist, however, asks us to discard the theoretical portion of our foundation. The challenge he faces is to demonstrate that an adequate pragmatist foundation for logic can, in fact, be developed. Put in our terminology, the question is whether a purely pragmatist approach provides sufficient friction to produce a genuine foundation. The failure of well-known attempts to develop a purely (or even a largely) pragmatic foundation for logic (e.g., the conventionalist attempt, discussed in Chapter 9, Section 2) points to serious obstacles to a satisfactory foundation of this kind.

B. *The Centrality of Logical Constants/Operators to a Theoretical Understanding of Logic.* Although, as we have noted in the beginning of the present chapter, it is possible, in principle, to approach the task of constructing a foundation for logic from different directions, it has yet to be shown that a systematic foundation for logic will not have to include a systematic characterization of logical constants/operators. Once again, the challenge is to actually construct such a foundation. As far as I know, no detailed proposal for a foundation of this kind exists today.

C. *The Undergeneration of Our Criterion.* This issue will be addressed in Section 10.10.

D. *The Overgeneration of Our Criterion.* This is the most common objection, and it has taken two forms. The first (McCarthy, Gómez-Torrente, and Hanson) appeals to intuitive linguistic considerations—in particular, considerations of what natural-language speakers would intuitively consider a genuine logical consequence. The second (Feferman) appeals to considerations coming from metamathematics and the philosophy of mathematics. I will briefly respond to the former and then offer a more detailed discussion of the latter.

1. *The overgeneration objection—I: Intuitive linguistic considerations.* The intuitive overgeneration objection says that our specialized criterion for logical constants overgenerates on the ground that some of the constants satisfying it give rise to consequences which are *intuitively* non-logical. For example, some people think that

⁷⁵ My discussion of (A) through (D1) will be relatively brief and will relate to generic views rather than to specific critics. My discussion of (D2) will be more detailed and specific.

the quantifier expression “the number of moons of Earth” satisfies our criterion, and this leads them to object to our criterion on the ground that it leads to the acceptance of consequences which are intuitively not logical (indeed, empirical) as logical, for example the consequence

(39) (The Number of Moons of Earth x) Φx ; therefore: $(\exists x)\Phi x$.

This objection is based on a misunderstanding of either our approach to natural-language expressions or our criterion of logical constants in Section 10.4, which says, among other things, that a logical constant must be an extensional rigid designator. Is the quantifier expression in (39) an extensional rigid designator? That depends on (a) how one understands “the number of moons of Earth” in (39), and (b) how one understands “an extensional rigid designator” in our criterion of logical constants. My claim is: words can be used in many ways, and their meaning may change from one use to another. It is not our job to legislate how people should and should not use such expressions as “the number of moons of Earth”. But it is our job to clarify under what uses natural-language expressions are extensional rigid designators. On my understanding of “an extensional rigid designator” in the present context, “the number of moons of Earth” is *not* an extensional rigid designator *when* it is used as it is commonly used in natural language, but it *is* an extensional rigid designator *when* it is used as *synonymous* with “there is exactly one”, where “one” is the canonical name of the second-level cardinality property ONE. What is a *canonical* name of ONE? In the present context, it is a name given to ONE by an adequate background theory of formal structure used by logic. Since such a theory, on our account, is a *mathematical* theory, the name in question must be a mathematical name, i.e., a name that does not use non-mathematical vocabulary like “moon of Earth”. “Exactly one” is a canonical name of ONE; “the number of moons of Earth” is not. But since we cannot tell people how to use “the number of moons of Earth”, all we can say is that (39) is a logical consequence on our account *if* its words are so used that it is synonymous with:

(40) (There is Exactly One x) Φx ; therefore: $(\exists x)\Phi x$

(where “exactly one” is used as it is commonly used), but *not if* its words are used in the way those who take their cue from the common usage use “the number of moons of Earth”.⁷⁶

2. *The overgeneration objection—II: Feferman’s metamathematical considerations.* Feferman (1999, 2010) raises three objections to our criterion for logical operators, which he calls “the Tarski-Sher thesis” (Feferman 1999: 37):

- (a) “The thesis assimilates logic to mathematics, more specifically to set theory.”
- (b) “The set-theoretical notions involved in explaining [the thesis] are not robust.”

⁷⁶ This is a new formulation of the explanations I gave in Sher (1991 Chapter 3, 2001, and 2003).

- (c) “No natural explanation is given by [the thesis] of what constitutes the *same* logical operation over arbitrary basic domains” (Feferman 1999: 37).

Let me begin by putting Feferman’s criticisms in perspective. There are many commonalities to Feferman’s approach to logic and ours. Feferman is involved in a serious foundational project,⁷⁷ he recognizes the need for a theoretical criterion for logical operators, he appreciates the philosophical significance of such a criterion, and he is not averse to a semantic formulation of a criterion of this kind. Furthermore, Feferman accepts the invariance-under-isomorphism criterion as a *necessary* condition on logicity, and he does not object to the view that logicity is grounded in formality. Finally, Feferman offers only a limited revision of this criterion. Nevertheless, Feferman is highly critical of the “Tarski-Sher” thesis, or of the claim that its criterion is sufficient for logicity. Let us examine his objections one by one.

(a) *The Tarski-Sher Thesis Assimilates Logic to Mathematics*. Feferman says:

The first of these [objections], I think, speaks for itself, . . . but it will evidently depend on one’s gut feelings about the nature of logic as to whether this is considered reasonable or not. For Sher, to take one example, this is no problem. Indeed, she avers that “the bounds of logic, on my view, are the bounds of mathematical reasoning. Any higher-order mathematical predicate or relation can function as a logical term, provided it is introduced in the right way into the syntactic-semantic apparatus of first-order logic.” ([Sher 1991], pp. xii–xiii) What that “right way” is for her, is spelled out in a series of syntactic/semantic conditions . . . ([Sher], pp. 54–55) . . . [One consequence is that] we can express the Continuum Hypothesis and many other substantial mathematical propositions as logically determinate statements on the Tarski-Sher thesis. . . . But in so far as . . . the thesis requires the existence of set theoretical entities of a special kind, or at least of their determinate properties, it is evident that we have thereby transcended logic as the arena of universal notions independent of “what there is” (Feferman 1999: 37–8).

First, I should say that (unlike Tarski) I agree with Feferman about the existence of significant, objective differences between logic and mathematics.⁷⁸ But the existence of such differences is compatible with the existence of systematic connections between the two. The invariance-under-isomorphisms criterion, as a criterion that grounds logicity in formality, draws a close connection between logic and mathematics. But it does not “assimilate” logic to mathematics or to set theory, in

⁷⁷ See Feferman (1984, 1993a,b) and Feferman and Hellman (2000).

⁷⁸ Tarski thinks that one’s attitude toward the assimilation of logic and mathematics is a subjective matter, a matter of a difference between two “types of mind” (Tarski 1966: 153)—the philosopher’s and the mathematician’s:

A monistic conception of logic, set theory, and mathematics . . . appeals, I think, to a fundamental tendency of modern philosophers. Mathematicians, on the other hand, would be disappointed to hear that mathematics, which they consider the highest discipline in the world, is a part of something so trivial as logic; and they therefore prefer a development of set theory in which set-theoretical notions are not logical notions (Tarski: 153).

particular when mathematics (set theory) is construed as a first-order discipline.⁷⁹ One manifestation of this is the fact that mathematical individuals and most of their properties (relations, functions) *do not* satisfy this criterion. Thus, no mathematical individual—neither 2 nor \aleph_0 nor any other cardinal or ordinal or set as individual—is logical,⁸⁰ and it is easy to see that the majority of mathematical properties of individuals—even, odd, prime, is a member of (\in), etc.—do not satisfy it either.⁸¹ So if by saying that our criterion “assimilates” logic to mathematics Feferman means that it “identifies” logic with mathematics or includes logic in first-order mathematics, his claim is inaccurate. On our account logic and mathematics stand in a systematic relationship to each other, but they are not *identical* or included one in the other.⁸²

Feferman’s first objection, however, raises other issues as well. One of these is the role of commonsense intuition, or “gut feelings”, in determining our attitude toward issues like the relation between logic and mathematics. Feferman says that “it will evidently depend on one’s **gut feelings** about the nature of logic as to whether [assimilating logic to mathematics] is considered reasonable or not”. I disagree. Everyday intuition might play a useful role in the first stage of a philosophical investigation, but later on its role should be limited, and theoretical considerations should take center stage. In fact, elsewhere Feferman himself describes the tasks of foundational studies as calling for *theoretical* rather than intuitive considerations. Their tasks include “*conceptual clarification; interpretation [and] reduction... of problematic concepts and principles;... organizational... foundations; and reflective expansion of concepts and principles*” (Feferman 1993a: 106). All these are theoretical rather than intuitive pursuits, pursuits that might very well end up in conflict, rather than in agreement, with our “gut feelings”.

Further issues raised in the above citation are *ontological commitment* and *substantiveness*. Feferman seems to uphold the traditional view that logic, unlike mathematics, should have no ontological commitments and make no substantive claims. The invariance-under-isomorphism criterion, Feferman says, burdens logic with both.

My response is, first, that one has to distinguish between the ontological commitments of logic as they arise from the general version of the invariance-under-isomorphisms criterion—logicality as formality in the general sense—and as they arise from its specialized version. The general version of our criterion of logicality

⁷⁹ In fact, our account does not assimilate logic to higher-order mathematics either.

⁸⁰ To be logical/formal is to abstract from the identity of individuals, hence no individual is logical/formal. Technically: the identity of individuals is not preserved under isomorphisms.

⁸¹ To see why the first-level membership relation (\in) is not logical, consider a structure in which it holds between two individuals, say a structure with the two individuals 1 and {1}: $\langle\{1, \{1\}\rangle, \langle 1, \{1\}\rangle$. This structure is isomorphic to the structure $\langle\{1, \text{Barack Obama}\rangle, \langle 1, \text{Barack Obama}\rangle$, but the membership relation is not preserved under this isomorphism: whereas 1 is a member of {1}, 1 is not a member of Barack Obama. (While the pair $\langle 1, \{1\}\rangle$ satisfies \in , the pair $\langle 1, \text{Barack Obama}\rangle$ does not satisfy \in .)

⁸² We will discuss this relationship in greater depth in Sections 10.7–8.

commits us to the reality of formal properties in general, but not to any specific formal properties. What formal properties there are is a question for other branches of our system of knowledge (in particular, our mathematical theory of formal structure) to answer. It is only when we choose a specific background theory of formal structure that our criterion—in its specialized version—involves logic in something like ontological commitment, and even then, commitment to *formal properties*, *formal laws*, and *formal possibilities*, rather than to *actually existent individuals*.⁸³

Second, the view that logic cannot have ontological commitments of any kind and the view that logical truths must be trivial have never received adequate justification. They are remnants of traditional foundationalism, Platonism, and conventionalism—problematic doctrines that we reject. Our examination of logic earlier in this chapter suggests that while logic does have special traits it is neither “pure” nor trivial. One of the points we have tried to make in this essay is that philosophy does not force us to choose between giving an *absolute* foundation to human knowledge and giving *no* foundation, between a conception of logic as *pure*, *rarefied*, *lofty* and a conception of logic as *trivial* or *empty*.

(b) The Set-Theoretical Notions Involved in Explaining the Tarski-Sher Thesis are Not Robust:

Point 2 is in a way subsidiary to point 1. The notion of “robustness” for set-theoretical concepts is vague, but the idea is that if logical notions are at all to be explicated set-theoretically, they should have the same meaning independent of the exact extent of the set-theoretical universe. For example, they should give equivalent results in the constructible sets and in forcing-generic extensions. Gödel’s well-known concept of absoluteness provides a necessary criterion for such notions and, when applied to [the kind of operators considered by the Tarski-Sher thesis] considerably restricts those that meet this test. For example, the quantifier “there exist uncountably many x ” would not be logical according to this restriction, since the property of being countable is not absolute (Feferman 1999: 38).

Feferman, however, qualifies his support of the absoluteness criterion somewhat:

One should be aware that the notion of absoluteness is itself relative and is sensitive to a background set theory, hence again to the question of what entities exist (Feferman: 38).

Thus, Feferman’s claim is that many of the logical operators sanctioned by the invariance-under-isomorphisms criterion are couched in set-theoretical notions that are not “robust”, and he argues that only operators couched in robust set-theoretical notions should be classified as logical. By a “robust” set-theoretic notion Feferman understands a notion or a term having “the same meaning independent of

⁸³ Technically, the definition of a model as having a non-empty universe commits logic to the *necessary existence* of at least one individual, but as logicians and philosophers of logic often note, this is merely a technical commitment of current logical systems, motivated largely by pragmatic considerations, hence a commitment we could easily dispense with.

the exact extent of the set-theoretical universe". He acknowledges that this "notion of 'robustness' for set-theoretical concepts is vague" (Feferman: 38), but he suggests that we partially identify it with the clearer (though non-robust) notion of "absoluteness".

The notion of *absoluteness* was introduced by Gödel in the course of proving the relative consistency of the Axiom of Choice and the Generalized Continuum Hypothesis (Gödel 1940), and it has been further studied by Barwise (1972), Burgess (1977), Väänänen (1985) and others. From the point of view of Feferman's criticism, the most relevant feature of the absoluteness requirement is that the meaning of a given characterization of a logical operator should not be affected by expansions or contractions of a given universe. Feferman defines absoluteness as follows: Given a set, T , of axioms in the standard language of set theory, a "formula φ of set theory is defined to be *absolute with respect to* T if φ is invariant under end-extensions for models of T " (Feferman 2010: 13). Which operators are absolute, i.e., defined by absolute set-theoretical notions in Feferman's sense? The existential and universal quantifier operators, for example, are absolute, but the *uncountably many* operator is not absolute. Are all non-standard quantifiers satisfying our criterion non-absolute? No. The *well-foundedness* quantifier, for example, is absolute.⁸⁴

The crucial question is whether the non-absoluteness of some of our operators is a good reason for rejecting our criterion. I believe it is not.

First, the relevance of robustness or absoluteness to *logicality*, let alone to the idea that logicality is grounded in formality, is highly questionable. Certainly, its relevance to logicality has never been demonstrated.

Second, the robustness objection is directed only at the specialized version of our criterion, not at its general, and philosophically more significant, version. The general version of our criterion is not inherently connected to any particular set-theoretical language, the kind of language for which the question of "absoluteness" might arise.

Third, even with respect to the specific version of the logicality criterion in terms of isomorphisms, a robustness criticism is problematic. For one thing, the notion of robustness, as Feferman acknowledges, is itself not robust. This raises the question: why, if non-robust notions should be avoided in a foundation of logic, should they be used in setting constraints on the descriptions of logical operators? Feferman does not address this question. Moreover, even if robustness/absoluteness is a desirable feature of theories in general, there are many other desirable features of theories that Feferman does not insist a logical theory should have. Take, for example, *decidability*. Decidability is a desirable property of theories, yet even logicians who limit themselves to the standard logical operators do not limit themselves to a logic that satisfies the decidability desideratum.

Fourth, not only does the robustness criticism concern an *artifact* of a particular background theory (one to which we are not wedded), but this artifact is largely due

⁸⁴ See e.g. Jech (2003: 502).

to a limitation of the logical framework used in formulating this background theory, a limitation that the very criterion targeted by the criticism enables us to overcome. Let me explain. The fact that many standard set-theoretical notions are not robust, in one sense or another, is connected to the fact that standard set theory is formulated within the framework of standard first-order logic. That is to say, to a significant extent, the non-robustness of many set-theoretical notions is due to the expressive limitations of the logical framework within which they are formulated, a logical framework whose logical operators (other than sentential connectives) were selected without appealing to any systematic principles.

One type of non-robustness of this kind is the one associated with the Löwenheim-Skolem (LS) phenomenon or with the phenomenon of non-standard models more generally.⁸⁵ These phenomena destabilize, or render non-robust, multiple notions defined within the *standard* first-order framework. The first-level notion “*x* is uncountable”, for example, as defined in standard set theory, is satisfied by *countable* sets in LS models, and this renders it “unstable” or “non-robust” in first-order set theory. It is easy to see that this problem is directly related to the limited expressive resources of the standard first-order logical framework, which is closely connected to its limited trove of logical constants. This problem can in principle be solved by introducing a systematic criterion of logicity like ours, which sanctions the addition of logical constants that significantly increase the expressive power of the standard first-order logical framework. For example, by adding the quantifier “uncountably many” (either as a primitive or as a defined logical quantifier), we will be able to express the claim that there are *uncountably* many sets in such a way that it will have *no* countable models.⁸⁶ This is due to the fact that as a logical constant, “uncountably many” is hard-wired into the structure of a logical system to which it belongs, but its non-logical counterpart (in the present case, the first-level predicate “uncountable” of standard first-order set theory) is not. The meaning of the logical “uncountable” is fixed across models, but that of the non-logical “uncountable” is not. Its meaning is *highly variable*; i.e., it varies from model to model. This is a direct result of the design of logical systems in accordance with their role in our system of knowledge, as explained in Section 10.4.

So, while we use the non-robust notions of standard first-order set theory in formulating our (specialized) criterion, the existence of this criterion enables us to significantly reduce the non-robustness of these notions. This is an example of our foundational-holistic methodology at work. We use a less than ideal logical framework (the standard first-order framework) to establish a criterion that enables us to

⁸⁵ As noted above, LS says that if a theory formulated within the framework of *standard* first-order logic has an infinite model, it has a countable model. For non-standard models, consider, for example, uncountable models of arithmetic.

⁸⁶ See e.g., Keisler’s (1970) system of first-order logic with this logical quantifier, where the LS phenomenon does not arise.

improve this very framework. The process unfolds, in a Neurathian manner, in stages: In stage 1 we select an available background theory formulated within an available logical framework, aware that the latter generates destabilizing phenomena like LS. In stage 2 we construct a (specialized) criterion of logicity using this background theory (with its logical framework). In stage 3 we use this criterion to construct a stronger logical framework and a stronger background theory based on it. And in stage 4 we reformulate our criterion in more stable terms, using the stronger background theory (with its stronger logical framework).⁸⁷

Fifth: Feferman, as we have noted above, does not express any objections to the general principle underlying the invariance-under-isomorphisms criterion, namely, the principle that logical operators represent formal properties. But if this principle is accepted, then absoluteness is *not* a reasonable constraint on logical operators. The reason is that absoluteness requires that logical operators be blind to some *formal* differences between universes (i.e., differences in cardinality or size), but on the view that logic is grounded in the formal structure of reality, formal differences cannot be willy-nilly neglected.

Sixth: Absoluteness might be viewed as a means for preserving the *standard* boundaries of logic. But absoluteness is not a suitable means for this goal. Absoluteness does not distinguish between standard and non-standard logical operators, as we have noted above. (Some non-standard logical operators are not absolute, some are.) Furthermore, the bias toward the standard, narrow notion of logicity is countered by the principle that, all things considered, it is usually preferable to study central notions in *full generality*⁸⁸ rather than limit their study in the hope of “avoiding problems”. In their book on the foundation of set theory, Fraenkel, Bar Hillel, and Levy express this principle as follows:

It may be safely stated that . . . throughout mathematics—and other disciplines—the investigation of the most general notions, **in all their unrestricted generality**, has often proved to be of extreme value for the advancement of research. To think that difficulties could be overcome simply by disregarding the general case . . . is somewhat naive (Fraenkel et al., 1958/73: 3).

Our interest in a *general* criterion of logicity is a reflection of this principle.

Finally, it is important to distinguish between a *criterion of logicity* and *choice of a logical system*. While a criterion of logicity tells us which operators, in principle, can be used as logical operators by an acceptable logical system, this is just one of the many considerations involved in choosing a particular logical system to work with in a particular context or for a particular purpose. What system we choose to work with is guided by a variety of considerations, and it is at this level that considerations like robustness find a more natural place than at the level of a general theoretical criterion

⁸⁷ I will discuss this type of process further in the next section.

⁸⁸ By “full generality” we mean, of course, “full generality, short of triviality”.

of logicity. (We will return to the relation between choosing a criterion of logicity and choosing a logical system shortly.)

(c) *The Tarski-Sher Thesis Provides No Natural Explanation of What Constitutes the Same Logical Operation Over Arbitrary Basic Domains:*

It seems to me there is a sense in which the usual operations of the first-order predicate calculus have the *same meaning* independent of the domain of individuals over which they are applied. This characteristic is *not* captured by invariance under bijections. As McGee puts it “the Tarski-Sher thesis does not require that there be any connections among the ways a logical operation acts on domains of different sizes. Thus, it would permit a logical connective which acts like disjunction when the size of the domain is an even successor cardinal, like conjunction when the size of the domain is an odd successor cardinal, and like a biconditional at limits” ([McGee 1996], p. 577). . . . I . . . believe that if there is to be an explication of the notion of a logical operation in semantical terms, it has to be one which shows how the way an operation behaves when applied over one domain M_0 connects naturally with how it behaves over any other domain M'_0 (Feferman 1999: 38–9).

I believe this criticism, too, is unwarranted. First, in systematizing a theory we are often forced to accept entities which, viewed from outside the theory, look strange, appear to lack internal unity, and seem to have no rhyme or reason. Such entities make good sense within the theory, however, where their internal unity and *raison d'être* rest on the principles of that theory. This is something all mathematicians are aware of, and Feferman himself (2000) points to many mathematical objects that appear, in his words, “monstrous” or “pathological” yet are, he rightly insists, perfectly legitimate.

Second, the particular sentential connective used by Feferman (following McGee) as an example of a deviant logical operator sanctioned by the “Tarski-Sher thesis” is, in fact, *not* a logical sentential connective on *my* version of the thesis (although it is a logical connective on Tarski’s as well as Lindström’s versions of the thesis). This is because my logicity criterion for sentential logic is the usual Boolean or truth-functional criterion, and this criterion does not sanction logical sentential operators that take into account things other than the truth value of their arguments; in particular, it does not sanction logical sentential operators that take into account things like universes of discourse and their features.⁸⁹

- ⁸⁹
- (i) The reason Tarski and Lindström did not separate their logicity conditions of predicative operators from those of sentential connectives is, I believe, that they were not concerned with the philosophical significance of the criterion as an answer (or part of an answer) to the philosophical question “What is logic?” (Tarski 1966: 145). Instead, they used technical considerations of conciseness to determine how to formulate (their versions of) the criterion.
 - (ii) I should note that I failed to clarify this point when McGee first made it in his 1996 paper. At the time McGee asked me to comment on a draft of his paper, but I did not comment on this point since in the context of McGee’s paper it seemed to me to be of secondary significance.
 - (iii) Although I offered separate criteria of logicity for sentential and predicative operators, I also pointed out (in Section 10.4) that each logical sentential operator is correlated with a logical predicative operator (e.g., negation with complementation) which, as such, may take the

Third, predicative logical operators which change their behavior from universe to universe are not special to our criterion. Everyone who accepts standard first-order logic as a legitimate logic accepts logical operators that seemingly lack internal unity. Consider a quantifier Q which behaves like “all” (\forall) in universes of fewer than 101 individuals, like “some” (\exists) in universes with 101–745 individuals, and like “none” ($\sim\exists$) in universes of all other cardinalities. This quantifier is definable from the primitive logical constants of standard first-order logic and as such must be accepted as a perfectly legitimate logical operator by all adherents of standard first-order logic. Indeed, even standard sentential logic has “monstrous” logical operators. Consider a 132-place truth-functional connective that behaves like a conjunction on rows with exactly 2, 101, 103, 104, or 120–130 T’s (i.e., rows in which 2, 101, ..., sentential variables are assigned the value “True”), like disjunction on rows with 3, 4, 5, 6, and 70–100 T’s, and like some highly irregular, randomly generated connective on all other rows. Does this connective have “the same meaning” in all rows of its truth table (correlates of models)? Is there a natural connection between the way it behaves in rows with 100 T’s, rows with 101 T’s and rows with 102 T’s? But this logical operator is sanctioned by all adherents of standard sentential logic, and for a good reason. What makes it “the *same operator* in all rows” is the criterion of logicity for sentential operators, i.e., the fact that it is a *truth functional* or a *Boolean operator*.

Although Feferman’s criticisms are unwarranted, they bring to the fore a question that many have asked: is invariance-under-isomorphisms, or formality more generally, a necessary-and-sufficient condition for logical constants/operators or only a necessary condition?

II. *Is the invariance criterion for logical constants/operators necessary or necessary-and-sufficient?*

Most philosophers who have written on the invariance-under-isomorphisms criterion regard it as at least a necessary condition on logical constants and operators.⁹⁰ The question arises whether this criterion is a sufficient condition as well. Many of those who have given a negative answer to this question subscribe to the “standard first-order logic thesis” which, as we noted in Section 10.6, has never been adequately justified (at least on theoretical philosophical grounds). From our perspective, the question whether the criterion is both sufficient and necessary, as it is commonly asked, is ambiguous. Two different versions of this question are:

underlying universe into account (and in the case of complementation actually does). But as I also noted, these correlations are not arbitrary. And it is of special interest in the present context that logical sentential operators which intuitively have a clear identity are correlated with predicative operators which preserve their intuitive identity from one universe to another. (For example, negation is correlated with a predicative operator which behaves like complementation in *all* universes, regardless of their size.)

⁹⁰ We will discuss exceptions to this view in Section 10.10.

- (a) Is the fact that a constant/operator satisfies this criterion sufficient for introducing it into the logical system (or any logical system) we work with?
- (b) Will any choice of logical operators satisfying this criterion give rise to a universal logical system whose consequences transmit truth from premises to conclusion with an especially strong modal force?

These two questions put the formality/invariance-under-isomorphism criterion of logicity in quite different lights: from the point of view of (a), the criterion is necessary but not sufficient; from the point of view of (b), it is both necessary and sufficient. Let me explain.

(a) *Choice of a Logical System to Work with.* Question (a) brings to the fore the fact that there are multiple considerations for choosing a logical system, and at different times, and for different purposes, we choose our logical system based on different considerations. Suppose, for example, that we are looking for a logical system that will serve as a framework for a mathematical theory of formal structure. Then, at least in the early stages of developing such a theory, we will naturally look for a minimal logical system, i.e., a system that includes a minimal number of formal operators as part of its apparatus. Since we want to *study* these operators and *discover* their laws, we want to minimize the amount of information given by the framework itself, prior to our study. Or suppose that we value completeness so much that we want to limit ourselves to complete logical systems. This will limit our choice of collections of logical constants (though not just to the standard collection, as noted in Section 10.6).

Furthermore, considerations pertaining to special features of our background theory might affect the logical systems we choose to work with. Suppose we choose ZFC as our background theory of formal structure. Given that the Continuum Hypothesis is *independent* in ZFC, i.e., that it is not determined by ZFC whether 2^{\aleph_0} is identical to \aleph_1 , we may wish to avoid using a logical system that has both $\exists! 2^{\aleph_0}$ (“There are exactly 2^{\aleph_0} ”) and $\exists! \aleph_1$ (“There are exactly \aleph_1 ”) as logical quantifiers.

On the other hand, if we wish to prevent the destabilizing effects of the LS phenomenon, which is connected to the limited expressive power of the standard collection of first-order logical constants, we may add non-standard logical constants like the quantifier “uncountably many”, a move that will spare us from this phenomenon. Or, if our interests are linguistic, we may wish to add logical quantifiers like the 1- and 2-place “Most”, as in “(Most¹ x)Φx” and “(Most² x)(Φx, Ψx)”.⁹¹

From the perspective of question (a), then, our invariance criterion sets a necessary, but not a necessary and sufficient, condition on logical consequence.

(b) *Satisfaction of the Strong Generality and Modal Force Requirements.* If, however, our goal is to identify the totality of logical constants/operators which, in

⁹¹ See, e.g., Sher (1991) and Peters and Westerståhl (2006).

principle, give rise to highly general and modally strong consequences, then our invariance criterion—or at least the *general* version of this criterion—has the status of a *necessary and sufficient* condition. Since our interest in logic, being foundational, focuses on question (b), we give priority to the view that the invariance-under-isomorphisms criterion sets a *necessary and sufficient* condition on logicity (subject to the kind of limitations noted in (a)). But from the point of view of our epistemic project as a whole, both approaches to the invariance criterion are legitimate.

Another way to see this issue is from the perspective of two general principles of theorizing, *maximality* and *minimality*.⁹² The principle of minimality is motivated by the goals of economy, simplicity, and particularity; the principle of maximality—by the goals of generality, exhaustiveness, and informative explanation. Our criterion of logicity is maximalist. Maximalist theories often involve extravagant ontologies. Clearly, the ontology of ZFC—our current maximalist theory of formal structure—is extravagant. But even the ontologies of number theory and sentential logic are extravagant in the sense that the majority of the objects postulated by these theories have no practical or even theoretical interest on their own. Nobody will ever contemplate, let alone find a practical use for, the majority of objects in the natural numbers series or in the “universe” of Boolean truth functions. (Recall the connective mentioned in our discussion of Feferman’s third objection.) Yet to curtail these ontologies is to curtail the explanatory power of the respective theories.⁹³ Were we to characterize the logical constants by enumerating two or three or even ten of them, their nature would remain a mystery to us; but by accepting a rich ontology of logical entities, we gain insight into their nature and the principles governing them. Our economical loss is our explanatory gain.

III. Generality

In discussing the generality of logic in Section 10.5 we pointed out that while, due to its formality, logic is highly general, it is not *absolutely* general.⁹⁴ The question arises whether logic should be absolutely general, or more specifically, whether a maximally general theory of inference would do better than *formal* logic in performing logic’s job.

Before answering this question let us clarify what greater generality means in the present context. Fortunately, we have conceptual resources for a precise characterization of this notion: the larger the number of transformations under which the distinctive notions of a given discipline are invariant, the greater its generality⁹⁵.

⁹² These principles are related to the principles of generality and particularity discussed in Chapter 7, Section 2. Maddy (1997, 2007) talks about the maximality-minimality duality in connection with set theory.

⁹³ Recall our explanation of the large ontology of mathematics in Chapter 8, Section 4.

⁹⁴ See fn. 61 in Section 10.5.

⁹⁵ (i) For the source of this idea see Tarski (1966). (ii) As before, “larger” is understood in terms of inclusion.

Thus, since logical (formal) notions are invariant under *all* isomorphism transformations while physical notions are invariant only under *some* isomorphism transformations, logic is more general than physics.

Based on this characterization we may say that to be more general than logical (formal) notions is to be invariant under more transformations than isomorphism transformations. Is it possible for a notion to be invariant under more transformations than isomorphism transformations?—Yes. Invariance under isomorphism transformations is invariance under transformations induced by functions of a special kind: functions from one domain of individuals to a second (possibly identical) domain of individuals which satisfy the conditions of being 1-1 and onto. Now suppose we weaken these special conditions. We then arrive at a stronger type of invariance.⁹⁶ In the extreme case, suppose we remove both conditions. We then arrive at invariance under all functions whatsoever from one domain of individuals to a second domain of individuals. What notions are characterized by this type of invariance? Clearly not all the logical notions on our conception (i.e., all formal notions) or even all the logical notions on the standard conception. For example, non-identity (\neq) is not preserved under transformations based on functions from sets with more than one individual to sets with one individual. Or consider a still more extreme case, arrived at by replacing functions with relations: invariance under transformations induced by any (partial) relations between individuals from two domains. Even non-emptiness (\exists) is not preserved under transformations of this kind. What notions are characterized by invariance under transformations of this kind? *Semantic-* or *ontological-type* notions, i.e., such notions as “is an individual”, “is a 1-place property of individuals”, etc. Should we replace *formal* logic by *ontological-types* logic? Clearly not. If we replace the formal-logical operators by ontological-type operators, we will get an extremely limited set of logical operators, hence an extremely limited collection of logical consequences and an extremely limited theory of inference. This theory’s ability to perform logic’s epistemic task will be far inferior to the ability of formal logic to perform this task. Generality is an important epistemic norm up to a certain level, but beyond that level it becomes a liability.⁹⁷ This is another reason for characterizing logic as formal. Formality endows logic with high generality, higher than that of most other fields of knowledge, without depriving it of the power to provide a rich method of inference, one that achieves its designated goal of significant expansion of human knowledge.⁹⁸

⁹⁶ This idea is formulated in Bonnay (2008). One example of a stronger invariance discussed by Bonnay is due to Feferman (1999, 2010); another is due to Bonnay himself. I will briefly address these in Section 10.9.

⁹⁷ Recall our discussion in Chapter 7, Section 2 of the need to balance generality and particularity in order to arrive at substantive theorizing. An example of an absurdly general theory of nature is one consisting of the single sentence “Nature is the way it is”.

⁹⁸ Other things that formality can and generality cannot do (or at least has not been shown to be able to do) are to ground logic in reality and to provide a unified foundation for logic and mathematics. (For the latter, see Section 10.8.) Furthermore, while invariance under isomorphisms (formality) preserves a

IV. *Necessity*

In Section 10.5 we saw how logic's modal force can be explained by its formality: due to its strong degree of invariance, logic does not distinguish between most situations, including counterfactual situations. As a result, it holds in a very broad array of counterfactual situations, one that properly includes the counterfactual domains of most other disciplines, hence has an especially strong modal force. As in the case of generality, the necessity of logic is not absolute, but this is not an obstacle. Three critical questions that have been, or might be, raised in connection with the necessity of logic are: (a) Doesn't grounding logic in reality render it contingent? (b) Could not the formal laws grounding logic be different from what they are? (c) Do not problems due to ZFC undermine the necessity of logic? The first two questions concern our own account; the third was raised in a different context, but we will examine it with respect to our approach.

(a) *Doesn't Grounding Logic in Reality Render it Contingent?* Our answer to this question is negative. The view that by grounding logic in reality we render it contingent is based on a narrow materialistic view of reality, one that regards all claims grounded in reality as (automatically) contingent.⁹⁹ This view is often associated with an either-or conception of philosophy according to which the only alternative to extreme empiricism is frictionless Platonism. This either-or view is one we have challenged in the present essay. Physical objects (including physical properties), we have pointed out, have both material and formal properties, and these properties all "reside" in the actual world. But formal properties have a stronger degree of invariance than physical properties, and accordingly the modal force of the laws governing them is greater as well. Those who replace *degrees of modal force* by a simple "necessary-contingent" dichotomy restrict our ability to draw fine-grained modal distinctions. But be that as it may, on any reasonable construal of this dichotomy, laws that have the vast counterfactual scope of logical laws belong on the "necessary" fork.

(b) *Could Not the Formal Laws Governing Reality be Different from What They Are?* Another critical question that someone might ask is: Nothing you said about formal laws seems to imply that these laws could not have been different than they actually are. For example: suppose the formal laws governing the world are bivalent. This does not seem to conflict with the supposition that the formal laws could have been trivalent. But if the bivalent formal laws governing reality could be trivalent,

basic notion of object, the invariances associated with greater generality do not. For example, invariance under functions which are not 1-1 does not preserve identity and difference of objects, while invariance under functions which are not onto allows some objects (objects in its range) to vanish altogether. Arguably, discourse limited to these levels of invariance has a notion of object that is too thin to be useful in understanding logic.

⁹⁹ See Thalos (2013). A similar view was addressed in Chapter 8, Section 4 with respect to mathematics.

does this not make these laws contingent? Does it not mean that there are counterfactual situations in which the formal laws are trivalent rather than bivalent?

To answer this question we need to consider more aspects of modality than we have so far, something that would take us far afield. But briefly, and limiting ourselves to considerations that directly pertain to the present outline, we may approach this question as follows. First, let us assume, for the sake of brevity, that there is a single reality, with its own set of formal laws (formal in the sense of our general characterization of formality). Then these laws, being formal, have an especially high degree of invariance, hence an especially broad counterfactual scope and, therefore, an especially strong modal force—stronger than that of most other laws, for example, physical laws. Now suppose it is possible (in some unspecified yet significant sense of possibility) that our world had different formal laws. Then, being formal, they, too, would have an especially strong degree of invariance, hence modal force—stronger than that of what would then be the physical, biological, and other laws governing reality. This feature of formal laws would not change. Our claim is that logic has a very high degree of necessity in just this sense, and that it is this kind of necessity that is relevant to logic's ability to perform its designated epistemic role.¹⁰⁰

(c) *Do Problems Due to ZFC Undermine the Necessity of Logic?* Every theory requires a background theory, and logic is not an exception (unless you are a foundationalist). Standard mathematical logic uses ZFC as its background theory, and in presenting a specialized realization of our general conception of logic we, too, chose ZFC as our background theory. But we have emphasized that the adequacy of ZFC, or of any other particular theory, to serve in this role is, and will always remain, an open question, and it is perfectly legitimate to replace ZFC by another theory. We have thus insisted that it is important to distinguish between the *general* philosophical principles governing logic and their *specific* implementations using ZFC.

Many contemporary philosophers, however, do not draw this distinction, and this has led them to make (or to be understood by others as making) *sweeping* statements about what logic can and cannot do or be. The basis for such statements has often been alleged weaknesses of the standard background theory of logic, ZFC, and especially its limitation to set-sized structures, which is interpreted as *inability* to

¹⁰⁰

- (i) Another way of expressing this is in terms of “formal multiverses”. Our account is compatible with the possibility of formal multiverses, i.e., clusters of “possible worlds”, each with its own formal laws. Each formal multiverse would include multiple physical multiverses, i.e., clusters of “possible worlds” each governed by different physical laws but all governed by the same formal laws, hence all sharing the same logic. In each such world this shared logic would guarantee the transmission of truth from sentences to sentences with great generality and considerable modal force (i.e. in a way that is not affected by the physical laws governing this world). It is this kind of necessity which is relevant to logic's ability to perform its designated epistemic role in our world.
- (ii) We still have to explain how our logical method applies to our theory of formal structure, given that formal laws have the same modal force as logical laws, as well as how it applies to our logical theory. We will explain this in the next section.

handle larger, class-sized structures.¹⁰¹ Among those who raised, responded to, or discussed questions or criticisms centered on this theme are Kreisel (1967), Boolos (1985), Sher (1991, Chapter 3), McGee (1992, 2004), and Field (2008, 2009a, 2009b). The main criticism is directed at the Tarskian model-theoretic test of logical truth/consequence, viewed as a specific test indelibly connected to ZFC. Two forms this criticism has taken are:

- (i) An adequate model-theoretic semantics must provide counterexamples for all false claims of logical truth or consequence. Suppose the only type of structure which provides such a counterexample for some claim of logical truth/consequence, *S*, is a *class-sized* structure.¹⁰² Then the Tarskian apparatus of models, being based on ZFC, will miss this counterexample. It will *wrongly* judge *S* to be logically true/valid.
- (ii) If a sentence is logically true then it must be true (simpliciter). Now, consider a set-theoretic sentence *S*. For this sentence to be logically true it must be true, i.e., true of all sets, of which there are class-, rather than set-, many. But the test of being true in all set-models does not include the case of truth in the class-size universe of all sets. Therefore the (ZFC) model-theoretic test fails to satisfy this requirement. One conclusion drawn from this consideration is: “[C]lassical models misrepresent reality: classical models have domains restricted in size whereas set-theoretic reality doesn’t” (Field 2009a: 264–5).

These and similar considerations have led philosophers such as Field (2008, 2009a, 2009b) to conclude that the conception of logical consequence as *necessarily truth-preserving* is misguided: “[L]ogic *can’t* be the science of what forms of inference necessarily preserve truth—even if the necessity in question is restricted to...*necessity by virtue of logical form*” (Field 2009a: 252). But if this claim is interpreted as a sweeping generalization, then it is not warranted. In particular, it is not warranted with respect to the foundational account of logic developed in the present essay. There is nothing in our foundational account (or even in our conception of the “Tarskian” approach to logical consequence) that requires the use of a background theory with a compulsory restriction on size. Most of what we say about logic is independent of ZFC.¹⁰³ Going from the general to the particular, we can divide our claims into five levels:

¹⁰¹ We briefly touched upon this issue in Section 10.4.

¹⁰² By “class” in this section I mean “proper class”. A proper class is a collection of objects which is too large or too unwieldy to be a set. To avoid Russell’s paradox, ZFC does not recognize proper classes as well-defined entities in the sense of entities one can quantify over. Significantly, the collection of all sets is a class.

¹⁰³ Mac Lane says a similar thing about the definition of categories. Starting with a “definition of a category *C* in a set-free form” (Mac Lane 1971: 24), he only later proceeds to reformulate the definition of *C* in set-theoretical terms, enabling us to understand *C* both in a general way and in a specific way, where only the latter is subject to the special restrictions introduced by set theory. (See discussion in Mac Lane 1971: 21–4.) This corresponds to our general and specialized characterizations of logicity, where the modal force of logic is first explained on the general level.

Level 1: Logical truths and consequences are significantly grounded in *reality* (as opposed to being grounded only in the mind).

Level 2: Logical truths and consequences are grounded in the *formal* facet of reality: specifically, in formal laws governing reality.

Level 3: Logical truths and consequences are truths and consequences that hold in all *formally possible structures* of objects (= scope of formal laws).

Level 4: Logical truths and consequences are truths and consequences that hold in all “*Tarskian models*”, in the *general sense of this term*, i.e., mathematical constructions that adequately represent the totality of formally possible states of affairs.

Level 5: Logical truths and consequences are truths and consequences that hold in all “*Tarskian models*” in a *narrow sense of this term*, i.e., models couched specifically in ZFC, hence set-models rather than class-models.

Field’s challenge applies to our account only on its fifth and most particular level.¹⁰⁴

Does Field’s criticism hold for our level 5 claim? Two challenges to a positive answer to this question come from (i) considerations concerning standard first-order logic, and (ii) considerations having to do with reflection principles of ZFC. The former appeal to the completeness of standard first-order logic together with the intuitive necessity of its axioms and rules of proof (Kreisel 1967; see also Sher 1991: 41–2). Field accepts these considerations and limits his criticism accordingly (Field 2008: 48). The latter appeal to “reflection” results saying that truth in the class-sized universe of sets is adequately represented by some very large set-structures.¹⁰⁵ Thus, McGee says:

[According to] *[r]eflection principles*... the universe of pure sets is so large and structurally variegated that every structural property of the universe as a whole is already exemplified at some ordinal level of the set-theoretic hierarchy (McGee 2004: 379).

Whether reflection principles are sufficient to allay Field’s worries is a question I will not go into here.¹⁰⁶

¹⁰⁴ (i) Another way in which Field’s context diverges from ours is in his greater interest in the use of logic in natural language. In contrast, we are primarily interested in its use in theoretical branches of knowledge.

(ii) It should be noted that the present discussion concerns only a small part of Field’s critical examination of logic, an examination that leads him to a wide array of positive as well as negative conclusions.

¹⁰⁵ For these results see e.g., Lévy (1960).

¹⁰⁶ I should note, though, that following the above citation McGee raises a question that might seem relevant to our specialized account of logical constants, namely, the logicity of “class” constants, i.e., expressions such as “absolutely everything”, construed as quantifiers [as in “(For absolutely every x) ϕx ”]. From our perspective, such quantifiers cannot be identified as logical by a specialized criterion couched in ZFC (though it can be identified as logical by a specialized criterion couched in a background theory that admits classes). It is, perhaps, worth adding that one can theorize about proper classes using a logical framework couched in ZFC, provided one treats class-expression as *non*-logical expressions.

10.8 Logic and Mathematics: An Alternative to Logicism

The great foundational systems of the late nineteenth and early twentieth century drew a close connection between logic and mathematics. Looking for a certain foundation for mathematics, *logicism* sought to ground mathematics in logic, *intuitionism* sought to ground both mathematics and logic in mental construction, and (*syntactic*) *formalism* sought to ground both in syntax. The present foundational inquiry diverges from the earlier ones both in its main goal and in its methodology. Our primary goal is a foundation for *logic* rather than for mathematics, and our methodology is *foundational holism* rather than the traditional foundationalism (with its emphasis on absolute certainty). But instead of breaking away from the logic-mathematics nexus, the new investigations have drawn us back to it, albeit with a new understanding.

Methodologically, unifying the foundations of logic and mathematics has advantages over distinct foundations. It enables us to replace two philosophical mysteries—the nature of logic and the nature of mathematics—by one. And instead of two difficult tasks—that of constructing a theoretical foundation for logic and that of constructing a theoretical foundation for mathematics—it presents us with one task. Unifying the foundations of two disciplines, X and Y, can take several forms. Among these are (i) reduction of Y to X, (ii) reduction of X to Y, and (iii) grounding both X and Y in a single principle, or a single cluster of principles, underlying both. A typical representative of (i) is logicism:

A. *Logicism*. Logicism is a well-known foundationalist program for grounding mathematics in logic, originating in Frege (1879, 1884, 1893) and connected to the nascence of modern logic. In addition to reducing two foundational tasks to one, logicism has the advantages of distinguished ancestry, familiarity, a rich body of literature, and important attempts at renovation (e.g., Hale and Wright 2001). But the foundationalist assumptions of logicism make it impossible for it to provide a foundation for *logic*. Logicism uses logic to explain and ground mathematics, but it lacks resources to explain and ground logic. Some philosophers have tried to solve the problem by suggesting a conventionalist foundation for logic, but this does not work since logical conventionalism itself, as we have seen in Chapter 9, Section 2, is highly problematic and, in particular, does not have adequate resources for a grounding of logic. An alternative to logicism falling under (ii) may be called “mathematism”:

B. *Mathematism*. Mathematism is as yet a largely unfulfilled option, although there are elements in intuitionism that can be viewed as partially realizing it (see e.g., van Dalen and van Atten 2002). While mathematism lacks logicism’s distinguished ancestry and rich body of literature, it shares the methodological advantage of reducing two foundational tasks to one. And since there exist several foundational accounts of

mathematics that do not put the main burden of explanation on logic—e.g., naturalism and structuralism—mathematism could, potentially, provide an adequate foundation for logic, provided it employs a holistic foundational methodology.

The third option is the one our own foundational inquiries lead to. Based on these investigations we will call it “formalism” or, to distinguish it from syntactic formalism, “semantic formalism”.¹⁰⁷

C. Semantic Formalism. Semantic formalism provides a common grounding for logic and mathematics in the formal facets of reality. Both logic and mathematics are grounded in *formal laws* governing reality, laws governing formal properties of, and formal operations on, objects in the world, though each in its own way. For example, the mathematical (set-theoretical) theorem “ $[t \cup (s - t)] = s$ ”, where s is a set and t is a subset of s , and the logical theorem “ $(\forall x)(Px \vee \sim Px)$ ”, where P is a 1-place predicate, are grounded in the same formal law. Theoretically, the formality of logic and mathematics consists in their discerning *formal* patterns, and only formal patterns, of objects and properties, i.e., patterns that do not distinguish the identity of individuals either within or across domains.

This joint foundation neither assimilates logic to mathematics nor assimilates mathematics to logic. Instead, it portrays logic and mathematics as standing in an intricate relation to each other, one involving a *common ground* on the one hand and a *division of labor* on the other. Both logic and mathematics are grounded in the *formal* aspect of reality, but they approach this aspect in different ways: mathematics *studies* the formal; logic *devises* a method of reasoning based on it. This difference is connected to another, related difference: the direct objects of logic are linguistic (logical truths and consequences), while the direct objects of mathematics are non-linguistic (numbers and sets, either as self-standing objects or as objectual posits). Furthermore, at least in the case of first-order mathematics, the objects of mathematics are of level 0 (individuals), while the objects of logic—and specifically their active ingredients, the logical constants and operators—are of levels higher than 0. Tarski, as we have noted earlier, thought it was a subjective matter whether one viewed mathematics as identical to logic or not. Our analysis points to a different conclusion: both the joint engagement of logic and mathematics with the formal and the division of labor between them with respect to it are *objective* matters, to be accepted or rejected based on impersonal considerations.¹⁰⁸

¹⁰⁷ Elsewhere I have also called it a *formal-structural* foundation and a *logico-mathematical structuralist* foundation. (See e.g., Sher 2001 and 2013).

¹⁰⁸ Our view of the relation between logic and mathematics focuses on mathematical fields such as arithmetic and set-theory, which are characterized by their subject matter and where mathematics is naturally viewed as teaching us something about the world, hence as subject to a genuine standard of truth. But arithmetic and set-theory are, of course, just two branches of mathematics. Mathematics is a diverse and multidimensional discipline, and mathematicians today are engaged in an enormous array of projects of multiple kinds. As a result, many would say that what distinguishes mathematics is its method rather than its subject matter. Our account explains the sense in which the mathematical method is *formal* as well:

We have talked about two components of the joint foundation of logic and mathematics: common grounding and division of labor. But there is a third component as well: cooperation. While logic uses mathematical knowledge as a basis for its inferential system, mathematics uses a logical inferential system as a framework for its theories. This takes us to our next topic, the intertwined development of logic and mathematics. Taking a functional perspective—one that focuses on the role logic and mathematics play in each other's development—we may describe their development as a cumulative process involving a continuous interplay between the two. Schematically, their interplay unfolds in a series of stages. Starting with a rudimentary logic-mathematical theory of formal structure, we develop a basic logic, which, in turn, produces resources for the development of a systematic theory of formal structure.¹⁰⁹ Then, using the resources of this theory, we develop a more advanced and sophisticated logic. And so on.

Painting with broad strokes, we can depict this series of stages as follows: arriving (in some early stage) at a rudimentary logic-mathematics that focuses on basic formal operations such as complementation, union, intersection, inclusion, and so on, we use these resources to generate a simple logical system (something on the order of syllogistic logic). Then, using this logical system, we develop a more sophisticated mathematics (say, up to and including nineteenth-century “naive” set-theory) as well as a more sophisticated logic (say, Fregean logic). Next, facing problems within the existent logic and mathematics (the paradoxes), we utilize resources developed in the last stage to construct a more attenuated logic (e.g., standard first-order logic with its standard proof system). And using this system as a framework for a more advanced mathematics, we develop rigorous and sophisticated axiomatizations of mathematical theories (for example, of arithmetic, Euclidean geometry, and Cantorian set theory). Some of these axiomatizations then lead to significant progress in our general theory of formal structure, say, ZF (and then ZFC). And using those, we are able to make considerable advancements in the metatheory of logic, including a systematic definition of logical consequence and the establishment of model theoretic semantics as an integral part of logic. This, together with other mathematical resources, enables us to develop a criterion of logicity (say, the invariance-under-isomorphism criterion). And that, in turn, opens the door for further developments in our theory of formal structure and, based on it, logic.¹¹⁰

a formal method is one that discerns only those features of the objects involved which are (either directly or indirectly) formal in our sense of invariance. (See also discussion of “mathematics as algebra” in Chapter 8, Section 4, where we point to another connection between so-called algebraic mathematics and veridical mathematics like set-theory or arithmetic.)

¹⁰⁹ Hanna (2006) and Maddy (2007) also think of logic as starting with something *rudimentary* (“protologic” for Hanna, “rudimentary logic” for Maddy), though they do not consider the possibility of a (functionally) intertwined development of logic and mathematics.

¹¹⁰ Many other examples branch out from this progression. One important example of this kind is the impact of model theory on mathematical developments concerning of infinitesimals (Robinson 1966/96).

This two-way, back-and-forth movement between logic and mathematics exemplifies what we have called “*constructive circularity*” in Chapter 2, Section 2. Among other things, circularity of this kind enhances our cognitive maneuverability in both fields. Russell’s paradox, for example, originally a problem for logic, was eventually overcome by relegating the problematic logical principle that gave rise to it (“comprehension”) to mathematics (set theory), where it was replaced by a mathematical axiom (“restricted comprehension”) that made it harmless. In that form it became available to logic through the use of axiomatic set theory as its background theory. This entire process did not render either mathematics or logic paradoxical. Nor did it render them trivial. The same logic can be used to develop diverse mathematical theories (including conflicting theories of formal structure), and the same mathematical theory can provide resources for developing diverse logics. This is what we meant in Chapter 2, Section 2 when we said that constructive circularity often involves “bringing new elements into the mix”. While each discipline provides crucial tools for developing the other, its own development crucially involves additional steps of discovery, invention, reflection, and reasoning that go beyond the other. The very act of *selecting* (*choosing, deciding on*) an appropriate background theory exemplifies this principle.

This holistic, back-and-forth pattern of interaction between logic and mathematics also explains how logic provides a modally powerful method of inference for the very theory that grounds it, the mathematical theory of formal structure. The fact that logic has an especially strong modal force is as significant here as in the case of, say, physics, though here the logical method does not have a *stronger* modal force than the theory it is applied to. This, however, is not a problem. Logic’s task is to develop a method of inference which is both highly general and modally powerful, and the generality of this method requires that it be applicable to all or most fields of knowledge, including fields whose own laws have a very strong modal force. This may not be compatible with foundationalism, but it is perfectly compatible with holism. Nevertheless, there is a clear sense in which it is logic that provides a method of inference for our theory of formal structure rather than the other way around. First, the theory of formal structure is not a theory of inference while logic is. Second, logic has an elaborate semantic-proof-theoretical apparatus that can be used to draw inferences in mathematics. And third, when we apply this apparatus to the theory of formal structure we hold the logical constants of this apparatus *fixed*, while when we study (within mathematics) the laws governing formal properties—including those denoted by our logical constants—we treat them as non-fixed. (For example, when we study *union* in set theory we treat union itself as not fixed in advance but disjunction—its logical correlate—as fixed in advance. If, following this study, we find out that our current disjunction is based on a faulty theory of *union*, we revise our logic in the next stage, using the holistic method.)

In a similar way we use logic in studying logic as well as in studying more general topics than logic, e.g., ontological types (is-an-individual, is-an-n-place-property-of-level-m, etc.).

Our discussion of logic and mathematics above adds another dimension to our account of mathematical truth in Chapter 8, Section 4. We began that account by pointing to the reality of formal features, and we said that mathematics' task (or one of its main tasks) is to study these features and the laws governing them. We have also said that although formal features are for the most part of second- and higher-level, mathematicians often use first-order theories to study them. And we have suggested that a possible explanation is that humans are better at studying certain (intricate) subject matters when they think about them, or represent them, in terms of individuals and their properties rather than in terms of properties of properties. All this retains its force here. But the present discussion suggests another reason for studying the formal by first- rather than higher-order theories: the first-order framework allows us more flexibility in determining what formal properties to hold fixed (as part of the logical framework) and what formal properties to leave unfixed, as non-logical elements whose laws are genuinely studied by the theory.¹¹¹ What mathematicians need is a logical framework that is neither so strong as to overpower their theories, nor so weak as to limit their ability to engage in a rich study of the formal. And it is easier to attain such a balance in first-order theories than in second-order theories.

But one result of using first-order theories to study the formal is that the objects these theories use to study it—say, the 0-level numbers of first-order arithmetic—are not themselves formal. They are non-formal posits representing formal objects. Now, this by itself is not a problem. We can learn about a huge metal-and-concrete building by studying a small plastic model of that building. But the question still arises where formality enters into first-order mathematical models. Our answer is that it enters through *structures*. The structures of the natural, whole, rational, and real numbers, unlike their 0-level members, are formal. They are invariant under isomorphisms. If N is a natural-numbers structure and N' is isomorphic to N , then N' is also a natural-numbers structure.

This answer points to a close connection between our conception of mathematics and that of the structuralists (Resnik 1981, 1982, and 1997, Parsons 1990, Shapiro 1997, and others). The connection is partly reflected in the centrality of *isomorphism* for both. To be formal, on our account, is not to distinguish between isomorphic structures, or to remain the same under isomorphisms, and this is also one of the main candidates for identity between mathematical structures according to structuralism.¹¹² This is expressed by Shapiro when he says:

¹¹¹ A third alternative is to hold some formal operators fixed as (second-level) operators of the logical framework and non-fixed as (first-level) operators of a non-logical first-order theory formulated within that framework. An example from first-order set theory is conjunction (in contexts of the form " $\Phi x \ \& \ \Psi x$ ") and intersection (as in " $x \cap y$ ").

¹¹² Note that I am not saying that the only way to think of mathematical structures is as invariant under isomorphisms. In some contexts, it is legitimate to think of, say, the structures $\langle N, 1 \rangle$ and $\langle N, 2 \rangle$, where 1

No matter how it is to be articulated, structuralism depends on a notion of two systems that exemplify the “same” structure. That is its point. . . . [W]e . . . need to articulate a relation among systems that amounts to “have the same structure”.

There are several relations that will do for this. . . . The first is *isomorphism*, a common (and respectable) mathematical notion. . . . Informally, it is sometimes said that *isomorphism* “preserves structure”. . . . [O]ne can informally take a *structure* to be an *isomorphism type* (Shapiro 1997: 90–2).¹¹³

A mathematical structure preserves its *mathematical identity* under isomorphisms just as a formal operator preserves its *identity* under isomorphisms. Both in the case of mathematical structures and in the case of formal operators we can say that identity is identity-up-to-isomorphism. There are many different structures of the natural numbers, and their identity as natural numbers structures does not depend on the identity of their individuals. Mathematical structures, like formal (logical) operators, do not distinguish the identity of individuals.

Indeed, even structuralists who do not identify isomorphism with structure identity regard it as central to structuralism. Thus Resnik views *isomorphism*, which he also calls “[p]attern-congruence” (Resnik 1997: 204), as the “strongest” among “the equivalence relationships which occur between patterns” (Resnik: 209). From a structuralist point of view there is no difference between studying the laws of arithmetic by studying one structure exhibiting the natural numbers pattern or another. For the *working* number theorist arithmetic may be a theory of a specific kind of individuals, but this does not conflict with the *philosophical* claim, reached by abstraction and generalization, that numbers are mere *places or positions in a structure*, where the identity of the “occupants” of these places is immaterial.

There are other points of similarity between our approach and those of mathematical structuralists. These include a realist outlook (e.g., Resnik 1997 and Shapiro 1997), a rejection of foundationalism and endorsement of holism (also Resnik and Shapiro), and a view of mathematical laws as having a strong modal force (e.g., Hellman 1989).

One difference is terminological. Whereas mathematical structuralists commonly speak about *structures* simpliciter, we speak about *formal structures*. In so doing, we do not just highlight the connection between mathematics and logic, but we also recognize the breadth of the notion of structure and its centrality to *all* theoretical disciplines. Different disciplines study structures of different kinds: physics—physical structures, biology—biological structures, and so forth. Mathematics studies formal structures. We may say that mathematical notions are formal in two ways: as representing higher-order formal properties (e.g., “1” representing ONE), and as

and 2 are individuals, as distinct structures. But mathematics *exhibits its formality* when it treats isomorphic structures as *interchangeable*.

¹¹³ Similarly, he says: “A purported implicit definition characterizes *at most one structure* if . . . any two models of it are *isomorphic* to each other” (Shapiro 1997:13).

representing places in structures ("1" representing the second-position in the ordering relation of the natural numbers structure).

Our conception of logic and mathematics as grounded in the same aspect of reality—the formal—enables us to utilize things we said about mathematics in our account of logic. In speaking about mathematical truth in Chapter 8, Section 4 we observed the ubiquity of the formal. Every object in the world has formal properties (e.g., self-identity); every property of objects in the world, or the extension of such a property, has formal features (e.g., cardinality); many relations exhibit formal patterns (e.g., reflexivity, symmetry, transitivity); and so on. This enables us to add another layer to our functional account of the emergence of logic in Section 10.1.

Thinking of the project of knowledge as targeting objects and systems of objects in the world, we note that objects and systems of objects have multiple properties: physical, biological, sociological, and many other kinds of properties, including formal or structural properties. One central task of theories of the world is to enable us to pass from existent knowledge to new knowledge. For example, if we know that water is the chemical compound H_2O , we can use our knowledge of chemical compounds to expand our knowledge of phenomena involving water. But while chemical, biological, sociological, and other kinds of laws contribute to our understanding of a limited array of phenomena, formal laws contribute to our understanding of almost every facet of the world. In every area of the world individuals are *identical* to (*different from*) individuals; objects lie in the *intersection*, *complement*, and *union* of properties; properties are *included* in properties; relations are (or are not) *transitive*, *reflexive*, *well-ordered*; properties and relations are *non-empty* or *universal*, *finite* or *infinite*, hold of n elements (tuples of elements) or m elements, and so on. Given the prevalence of formal properties and relations, the possibility of developing a wholesale method for extending knowledge based on universal laws governing the formal properties of objects arises. Such a method will enable us to move from existent beliefs to new beliefs without loss of truth and regardless of field of interest. The task of logic, on this approach, is to construct a theory of the transmission (preservation) of truth based on formal or structural grounds. The formal laws underlying this theory are studied by mathematics; their formal parameters are built into our language as logical constants. This, in turn, gives rise to "logical forms", and these are then connected by a method of inference based on, or applying, the laws studied by mathematics, which, being formal, have an especially strong modal force. While all formal laws are in principle universally applicable, some are used more widely than others, or more widely in certain contexts. For example, the laws of union and intersection have multiple applications in all areas of knowledge, but the laws of infinite cardinalities do not. This partly explains why the standard logical system includes logical operators based on union and intersection but not on infinite cardinalities. In principle, however, laws governing the latter can ground consequences with the same modal force and generality as laws governing the former.

10.9 On the Possibility of Error and Revision in Logic

Logic, like any other branch of knowledge, is susceptible to error. Its theories, like those of other disciplines, are subject to significant veridicality standards, and their success in satisfying these standards cannot be taken for granted or be guaranteed prior to a substantive critical examination. This is a significant aspect of our claim that logic lies in the periphery. A paradigmatic example of error in logic, one we have noted earlier in this outline, is the error discovered by Russell in Frege's ground-breaking logical theory. Many philosophers, however, view logic as immune to error. Some of the (explicit or implicit) arguments in support of this view say that (a) logic is analytic, (b) logic is obvious, (c) logic is not subject to empirical tests, and (d) logic is necessary.

(a) *Analyticity*. One prevalent view is that since logic is analytic rather than factual, there is no question of (factual) error in logic. Change in logic is nothing more than a change in meaning; there is no theoretical change in logic. We have explained why we reject the view that logic is analytic. Field's argument against the view that change in logic is merely a change in meaning captures one aspect of our position. Targeting Carnap, Field says:

Carnap...thought...the logical connectives differ in meaning when used by advocates of different all-purpose logics. But the notion of difference of meaning is unhelpful in this context. On some readings of "differ in meaning", any big difference in theory generates a difference in meaning. On such readings, the connectives do indeed differ in meaning between advocates of different all-purpose logics, just as 'electron' differs in meaning between Thomson's theory and Rutherford's; but Rutherford's theory disagrees with Thomson's despite this difference in meaning, and it is unclear why we shouldn't say the same thing about alternative all-purpose logics. Of course if the connectives differ in meaning between the theories in a more substantial sense—if, for instance, the proponent of one all-purpose logic should translate the proponent of the other non-homophonically—then that might suffice to remove disagreement between the two. But it is hard to see what grounds there could be for a claim of difference of meaning in any such strong sense, or indeed, in any stronger sense than that exemplified by the difference of meaning in 'electron' from one theory to another (Field 2009b: 345).

(b) *Obviousness*. Here the thought is that logic is too obvious to be fallible. We have argued against this view (as a vestige of foundationalism) in Chapter 9, Section 2, so there is no need to elaborate further.

(c) *Empiricism*. An (often implicit) argument for the infallibility of logic is associated with radical empiricism. Since logic is immune to empirical refutation, so the reasoning goes, it is immune to any kind of factual refutation. The idea is, roughly, that since the only measure of (factual) error is experience and logic is not empirical, logic must be immune to error. We reject (on general epistemic grounds)

the identification of facts with empirical facts and the sharp bifurcation of knowledge into empirical and non-empirical inherent in this reasoning. And while we do acknowledge logic's immunity to most empirical discoveries, we explain this in a way that does not imply it is immune to all factual discoveries, or even to those having some empirical component. On the contrary, we identify a type of fact that logic is highly sensitive to, namely, facts concerning formal laws governing reality. Logic is affected by discoveries concerning such facts, whether they have any empirical element or not.

(d) *Necessity*. Another "reason" for viewing logic as infallible is its strong modal force. Since logical truths are *necessarily true*, one might reason, they *cannot be false*; but if they cannot be false, they are *infallible*. This line of reasoning is based on a confusion. The formal necessity of the logical laws does not imply that logic is infallible any more than the nomic necessity of the physical laws implies that physics is infallible. Newtonian and Einsteinian laws differ (as far as we know) in their truth, but they do not differ in their modal force. We can be wrong about what the laws of nature are without being wrong about their modal status. And the same holds for logical laws. One can doubt the logical truth of the classical law of excluded middle without doubting that if it is logically true it is necessarily true. To say that the logical truths are necessary is to say that *if* a given sentence is *in fact* logically true, then it is also (in fact) necessarily true; it is not to say that all the sentences that a given logical theory *regards* as logically true are *in fact* necessarily true, hence *true*. We conclude that contrary to the views considered above, our logical theories are fallible.

What, then, are some of the possible sources of error in logic? One potential source of error in logic was noted above: its background theory of formal structure. If the laws governing the formal features of objects (properties, relations, structures of objects, situations) are different from what logic's current background theory says they are, this is likely to result in errors in logic. For example, if the basic structure of properties is not bivalent, then the law of excluded middle is not logically true. Another potential source of error is the choice of logical constants. If we select "is taller than" or "is a property of humans" as logical constants (keeping the model theoretic definition of logical consequence unchanged), we will mistake material consequences for logical ones, and if we deselect the existential and universal quantifiers as logical constants, we will mistake logical consequences for material ones. A third source of error might lie in the construction of models. If we construct models as ranging (only) over *physically possible* structures of objects (instead of *formally possible* structures of objects), we will mistake physical laws for logical laws. And so on. Errors such as these may result in incorrect judgments of logical consequence, logical truth, logical consistency, etc., and as such may provide sound reasons for revision in logic.

Turning to revision, we view revision in logic, like revision in any other discipline, as warranted under a variety of conditions and based on a variety of considerations,

including factual (or periphery) considerations, methodological considerations, and pragmatic considerations, either separately or in various combinations. Factual considerations are considerations of truth and evidence, and they include (relevant) discoveries in the background theory of formal structure, both positive and negative. The latter might justify retraction and replacements of our current logical theory, the former, extensions of this theory. For example, Wiles's recent proof of Fermat's Last Theorem sanctions the introduction of the logical rule schema:

$$(41) \quad (\exists!k^n x)\Phi, (\exists!l^m x)\Psi, \sim(\exists x)(\Phi \ \& \ \Psi) \vdash \sim(\exists!m^n x)(\Phi \vee \Psi),$$

where $n > 2$ and $k, l, m > 0$.

One common objection to revision in logic says:

[T]here are certain principles of logic . . . that cannot be given up, because they are framework principles . . . What would happen if we gave up the law of non-contradiction? It is not clear that there could any longer be such a thing as what Quine calls a recalcitrant experience, forcing changes elsewhere in the system. If a new deliverance stood in contradiction to things we already accepted, we could simply accept it with a "What—me worry?" shrug (Van Cleve 2005: 177).

In response to this objection, let me say first that our bedrock principle, which we share with Van Cleve, is that a viable system of knowledge requires a viable logic, and that a viable logic must play a substantial role in constraining our system of knowledge in a variety of ways, for example, making sense of, and giving teeth to, the idea of "recalcitrant experience" or more generally "negative evidence". This, however, does not mean that a particular logical theory—say, the logical theory accepted today—or a specific theorem of this theory—say, the law of non-contradiction—is required for logic to perform this role. Second, it is important to note that revision in logic, as we envision it here, is not revision of a single law in a vacuum. More likely, it is a (partial) replacement of one logical system by another, so the loss of a principle like the law of non-contradiction is compensated by the introduction of other principles that play the same or a similar role. Third, the revision of a logical system, on our conception, is a serious matter, and any revision of our current logical system must satisfy stringent friction requirements, including the requirement of detecting formal conflicts between predictions and experience. This, however, does not mean that there is just one way of meeting this requirement. It is open for logic to meet this requirement in a variety of ways (and this openness is not conditional on anyone's ability to *imagine* today what an alternative to our current way of meeting this requirement might look like).

Let us conclude our present discussion of error and revision in logic with two questions. How can experience play any role in revision in logic, given that due to its strong degree of invariance logic is blind to experiential differences? What role, if any, can experience play in discovering error and initiating revision in logic? In his letter to Morton White (see Section 10.2) Tarski says:

Axioms of logic are of so general a nature that they are rarely affected by . . . experiences in special domains. However, . . . I can imagine that certain new experiences of a **very fundamental nature** may make us inclined to change just some axioms of logic. And certain new developments in quantum mechanics seem clearly to indicate this possibility (Tarski 1944a: 31–2).

One has to be careful in attributing to an author a view informally delineated in a personal letter, so I will speak only for myself. Due to the strong invariance of the formal, formal laws are too general, too abstract, and too modally strong to be directly discoverable by empirical methods. But the possibility that some combination of experiential and theoretical considerations will reveal an error in logic or suggest the superiority of one theory of formal structure over another, and through this, of one logic over another, cannot be ruled out. There is no question here of a single empirical experience (or even a series of repeated experiences) that directly demonstrates the need to replace one logical theorem by another or one logical theory by another. But we cannot rule out the possibility of problems or discoveries concerning empirical science *pointing beyond* themselves to more abstract problems or discoveries, including ones pertaining to logic. Indeed, some adherents of quantum logic view the development of their logic in this way.¹¹⁴ Whether they are right or wrong, it is *in principle* possible that our understanding of the basic formal laws governing properties of objects in the world, and with them, the patterns of logical inference that work in the world, will undergo a change under similar circumstances.¹¹⁵

There is also a more direct way in which logic might be challenged by empirical discoveries, namely, challenges by empirical counterexamples to general logical laws. If, say, science tells us that neither *S* nor its negation is true, this poses a potential challenge to the logical law of the excluded middle. And if science tells us that *S* is true and *S'* is false, this is a potential counterexample to any logical theory which says that *S'* follows logically from *S*. Of course, these are also potential counterexamples to current science, but that does not rule out their being potential counterexamples to logic as well. Such counterexamples would be defeasible, yet under certain circumstances they could pose a genuine challenge to logic all the same.

It is important to emphasize that allowing experience a (limited) role in the revision of logic does not, by itself, render logic contingent (in the sense that the natural and social sciences are contingent). Nor does it interfere with its strong invariance. If and when revision is required, we replace our current logical theory by another with an equally strong modal force and equally invariant logical operators.

¹¹⁴ See the brief discussion and references in Chapters 3–4. For a judicious discussion of this case in the context of a foundation for logic see Maddy (2007, Part III).

¹¹⁵ Here is a domestic example of how one set of problems can point to another: A couple sits down at their dining table, having dinner. One spouse complains about the food prepared by the other. This develops into a fight about various mundane issues of their daily life. At some point the two realize that these mundane problems point to an altogether different and deeper problem, one they have never suspected and that calls for a major change, on an altogether different level, in their marital life.

What about justification? Clearly, on a holistic view of justification, there is no deep chasm between justification based on purely intellectual considerations and justification based on a combination of intellectual and empirical considerations. And due to substantial interconnections between different fields of knowledge, considerations coming from different fields—e.g., abstract and empirical fields—can contribute to theoretical justification in any field.

Another possible source of change in logic is change in other disciplines (or parts of disciplines) which are more general than, or as general as, logic. For example, changes in our basic view of objects might affect our conception of formal structure and with it logic. Finally, methodological and pragmatic considerations may lead to, or play a role in, revision in logic. Such considerations can be more general (e.g., the overall unity of our system of knowledge) or more specific (e.g., proof theoretic completeness). Generally, some combination of factual, methodological, and/or pragmatic considerations is involved in revision in any field, but veridicality considerations, on our conception, have a greater weight than, say, pragmatic considerations.

Needless to say, the revisability of logic extends to philosophical theories of logic, including the present theory. To date, the main proposals for revision in our theory have concerned its specialized criterion of logicity. Feferman (1999, 2010), for example, proposed a replacement of invariance-under-*isomorphisms* by invariance-under-*homomorphisms* as a criterion of logicity.¹¹⁶ Additional proposals were made by Bonnay (2008) and others. These proposals call for serious consideration, but so long as no attempt has been made to show that they add to, rather than detract from, our ability to provide a substantive philosophical foundation for logic,¹¹⁷ it is difficult to assess their merits.

10.10 The Scope of Logic

The scope of logic, on our account, is the scope of formal laws. Since formal laws apply, in principle, to all areas of knowledge, so does logic. Logic itself is a family of systems, each associated with some formal parameters, built into it as operators denoted by logical constants (primitive or defined). Which system we use on a particular occasion depends on our interests and goals on that occasion. But from the point of view of veridicality, generality, and modal force, all systems do, in principle, satisfy the standards set by logic's job description: the development of a highly general and modally strong method for transmitting truth from sentences to sentences (logical inference), the discovery of pernicious errors (logical contradictions and inconsistencies), and so on.

¹¹⁶ For discussions of Feferman's proposal see Bonnay (2008) and Sher (2008). Recently, however, Feferman has forsaken this proposal, replacing it with a new approach to the question of logicity, one which offers a "combined semantical and inferential criterion for logicity" (Feferman 2015). This proposal focuses on other issues than those discussed here.

¹¹⁷ Detract by, e.g., severing the philosophically significant connection between logicity and formality.

This conception gives rise to questions of scope of various kinds, some internal, others external. The former concern issues such as the use of logic in non-veridical contexts, e.g., command contexts or artistic fiction. The latter concern issues such as the status of modal logic and logical pluralism. Let us turn to these issues.

I. Non-veridical discourse

Our foundation for logic largely centers on veridicality concerns. Does this limit logic to veridical contexts? Consider the context of *commands*. Does the command:

(42) Answer all the questions on the exam!

logically imply the command,

(43) Answer question one on the exam!?

Our answer is positive: (42) does logically imply (43). But, someone might ask, what is the basis for this logical inference on your account, given that no transmission of truth is involved? The key to understanding why and how logic applies to, say, commands is the observation that commands, like veridical assertions, are *world-oriented*. Commands are given to agents in the world, refer to objects in the world, and concern acting in the world. Since the world is (objects and activities in the world are) governed by formal laws, formal laws are relevant to commands just as they are relevant to veridical assertions. In particular, they determine what is formally included in a particular command (answering all the questions on the exam includes answering the first question on the exam), whether a particular command is formally satisfiable (a command to answer all the questions on the exam and not to answer the first question on the exam is unsatisfiable), and so on. This explains why inferences based on such laws work in commands. Indeed, the inference from the command (42) to the command (43) is based on the same formal law as the inference from the assertion:

(44) John answered all the questions on the exam

to the assertion:

(45) John answered question one on the exam.

The same formal *law* that grounds the transmission of truth from (44) to (45) grounds the transmission of command from (42) to (43).

A more general perspective, applicable to a wider range of contexts, is the following: in any context whatsoever, veridical or non-veridical, we have to parse our subject matter in some way. Now, given that the same objects and properties are involved in both veridical and non-veridical contexts, and given that in the latter we have more freedom (we are less constrained) in parsing our subject matter,¹¹⁸ it

¹¹⁸ In veridical contexts we have to take into account considerations of truth, which are highly constraining, but in non-veridical contexts we are free of such considerations.

stands to reason that we give priority to the needs of veridical contexts in choosing a parsing for our entire language. Following this guideline, we construe (42) and (43) as having the same logical parameters as their veridical correlates, governed by the same formal laws. Consequently, the inference from (42) to (43) comes out logically valid, and its validity is grounded in the same laws as those grounding the inference from (44) to (45). That is to say, our method of inference for veridical contexts is naturally extendable to other contexts.

But does it extend to all contexts? In the extreme case: does it extend to highly fictional artistic contexts? Artistic contexts are contexts in which on the one hand the artist has considerable artistic freedom, but on the other hand she might still talk about objects and their properties, aim to communicate her ideas to her audience, and intend to say things that are relevant to life outside art. In such contexts, therefore, we find a mixture of logical and non-logical thought, and it is left to the author and reader to determine where to appeal, and where to suspend the appeal, to formal, hence logical, laws.

II. *Additional logics*

The scope of logic, on our conception, is in principle far broader than that of the currently prevalent paradigm of a logical system, namely, standard first-order mathematical logic. But there are many logical systems that extend this paradigm in ways that diverge from our extension. Modal, temporal, epistemic, deontic, and other logics extend standard first-order logic to types of reasoning that our conception does not cover. Often such theories are generated from standard first-order logic, or from sentential logic, by adding “fixed” or “distinguished” constants that do not satisfy our criterion of logicity, and they extend the standard semantics and proof theory to encompass forms of reasoning involving these constants. The new distinguished constants enable the theory to study inference forms involving such terms as “necessarily” or “believes”, and the extended theories often focus on the interaction between the new distinguished constants and those of standard logic, giving rise to such sentence forms as “Necessarily not S”, “Someone believes that $a=b$ ”, etc. These extended theories usually affirm all the inferences sanctioned by standard logic and add to these other valid inferences involving the new logical constants.

How are these logics viewed from our perspective—both from the perspective of our overall conception of knowledge and from the perspective of our foundational account of logic?

Theories extending standard first-order logic in the ways just noted fall squarely within the scope of knowledge on our conception. Our system of knowledge is a compendium of theories, each studying some aspect of reality, and together aiming at knowledge of all significant aspects of reality. The phenomenon of truth transmitting inference is one such aspect. This phenomenon encompasses multiple kinds of inference, and as such calls for multiple theories of inference. It is true that the inferences studied by these theories are in a certain sense weaker than those studied

by the logic our foundational outline focuses on. In particular, they have a weaker degree of invariance. But this does not impugn them as theories. On the contrary, both from the perspective of our general conception of knowledge and from the perspective of our conception of logic's role in knowledge there is a seamless continuity between theories falling under our conception and these theories.

Modal and other theories of inference are usually viewed as, and called, "logics" by their proponents. What are they from our perspective: Logics? Theories of extra-logical inference? Something else? First, both historically and thematically, there is much in common between these theories and the logical theories that fall under our conception—including goals, methods, and practices—and this justifies grouping them together within the same category. Second, it is built into our conception of knowledge that concepts such as "logic" are multidimensional and sensitive to perspective (see Chapter 4). This further strengthens our view that these theories fall under the category "logic".

Still, our decision to focus on a particular type of logic in this essay was neither arbitrary nor accidental. One cluster of reasons for this choice was historical and methodological. We were interested in a logic that had historically been considered one of the cornerstones of our system of knowledge yet was viewed as resistant to a foundation. Since the closest contemporary descendant of this kind of logic is the current system of mathematical logic, we chose to focus on this logic.

A more pertinent reason for focusing on mathematical logic was our conception of knowledge as requiring a method of inference that is (i) robustly veridical (grounded in reality), (ii) highly general, and (iii) modally strong to an exceptionally high degree. Since our investigations suggested that the key to possessing these three features was formality, the resulting logic turned out to be mathematical, or more precisely, a generalized version of standard mathematical logic. Thus, while our conception of knowledge encompasses logics of many kinds, mathematical logic is especially relevant to our project.

III. *Logical pluralism*

Another perspective on the scope of logic is offered by logical pluralism. Logical pluralists—Carnap (1934, 1939, 1950), Beall and Restall (2006), and others—claim that there is a multiplicity of *worthwhile* logical systems, including *conflicting* logical systems, such as classical logic and paraconsistent logic or intuitionistic logic. A characteristic motivation for logical pluralism is an attitude of *tolerance* (Carnap), *charity* (Beall and Restall), and *open-mindedness* (both). But while these are all worthy attitudes, a critical evaluation of logical pluralism requires a veridical perspective as well (indeed, first and foremost a veridical perspective). From this perspective, we may divide the space of logical pluralism to three positions:

- (a) There is no question of truth or correctness in logic, only a question of convenience. Different logics may be convenient for different purposes.

- (b) There is a question of truth or correctness in logic, but there are many incompatible true logics. Two conflicting logics can both be true.
- (c) Independently of the question of truth or correctness in logic, the development of new logics is a worthwhile enterprise, encouraging innovation and potentially leading to new discoveries.

Something along the lines of (a) is found in Carnap, (b) in Beall and Restall, and (c) in both.

Our approach to (a) and (c) is straightforward: we reject (a) and support (c). Carnap famously said: “*In logic, there are no morals*” (Carnap 1934: 52); “let any postulates and any rules of inference be chosen arbitrarily”¹¹⁹ (Carnap: xv); “[i]f we call [logical theorems] true, then another kind of truth is meant, one not dependent upon facts” (Carnap 1939: 2). This cluster of views we reject.¹²⁰ In logic, as in all fields of knowledge, there *are* morals, and one paramount moral logic shares with all fields of knowledge is “*Search for truth and avoid error*”. Knowledge, qua knowledge, requires a standard of truth, and logic, as a field of knowledge, is not exempt. Whatever other constraints logic is subject to, its judgments of logical consequence are constrained by the need to be true. This means that we cannot choose our logical constants or axioms or rules of inference arbitrarily: an arbitrary choice is at least as likely to result in false judgments of logical consequence as in true judgments. Logic, indeed, is required to be true—true to reality—on two levels. First, it is required to be true to the reality (facts) of what follows logically from what, namely: which sentences transmit truth to which sentences with an especially strong modal force, based on their logical form. This requirement is not trivial. If an inference fails to transmit truth from premises to conclusion with the requisite modal force, logic cannot arbitrarily say it succeeds. And second, whether an inference succeeds in transmitting truth from sentences to sentences depends on the connection—in the world—between the conditions that have to be satisfied (in the world) for its premises and conclusion to be true. This means that logic has to ground its inferences in *something in reality*, something that on the one hand plays a central role in the truth conditions of the sentences involved, and on the other hand has the requisite generality and modal force. In light of these non-trivial constraints on the choice (construction) of a logical theory, the acceptance of such a theory is far from a matter of arbitrary, or even purely pragmatic, choice.¹²¹

¹¹⁹ While it is conceivable that here Carnap had in mind something like the “mathematics as algebra” approach discussed in Chapter 8, Section 4, I use this citation to represent a more deeply conventionalist view.

¹²⁰ Rejecting the views expressed in these citations, however, does not necessarily mean rejecting Carnap’s view of logic in all its complexity. Our goal here is to address a certain line of thought that has become associated with some of Carnap’s pronouncements, including the ones cited above, but it is an open question whether, and to what extent, these citations are representative of Carnap’s overall view. (In this connection, see our discussion of Carnap’s conventionalism in Chapter 9, Section 2.)

¹²¹ Among pluralists who share our rejection of (a) are Beall and Restall (2006).

In contrast to (a), our foundational account of logic, and the more general account of knowledge in which it is embedded, support (c). One central task of our system of knowledge is to prepare tools for a variety of uses, leading to the charting of new domains of study, the development of new theories, and the improvement of existing theories. A paradigmatic example of (c) outside logic is non-Euclidean geometry in its early stages. Early non-Euclidean geometries satisfied our intellectual curiosity, addressed some questions of truth (e.g., the truth of the claim that the parallel axiom is independent of the other Euclidean axioms), and played an invaluable role in the development of truth-centered theories outside mathematics (especially, general relativity theory). Type (c) logical theories have similar virtues. They are expressions of epistemic freedom, opening a “boundless ocean of unlimited possibilities” (Carnap 1934: xv).¹²²

It is worth emphasizing that there is no conflict between our attitudes to (a) and (c). Our system of knowledge is a multifaceted, dynamic system, and at any given time there is room in it both for veridical theories, which are subject to a correspondence standard of truth, and for theories that fall under the category of “mere model” in the everyday sense of this word (as in our discussion of “mathematics as algebra” in Chapter 8, Section 4). Commitment to truth and an open-minded interest in models (which might eventually enable us to discover new truths) do not contradict one another.

But the true test of logical pluralism concerns (b), the view that there are many true or correct logics and that two conflicting logics can both be true. Today, it is hard to engage in a substantive discussion with logical pluralists on this issue since most pluralists avoid substantive foundational questions concerning logic, opting for a deflationist or a minimalist approach instead. Clearly, the less you say about logic, the more deflationist your outlook is, the fewer the constraints you set on the acceptance of any logic, and the easier it becomes to let the critical stance, so crucial for any field of knowledge, slide into an “anything goes” attitude.

The antidote to this attitude is a substantivist, foundational stance. Beall and Restall, for example, do not go as far as endorsing an “anything goes” methodology. They adhere to certain definite views concerning logic. For example:

We hold that deductive validity [or logical consequence] is a matter of the preservation of truth in all cases. An argument is valid when there is no counterexample to it: that is, there is no case in which the premises are true and in which the conclusion is not true (Beall and Restall 2006: 23).

However, they do not attempt to go beyond this minimalist characterization of logical consequence. They observe that there is no “settled” view about what the *totality of cases* relevant to logical consequence is, and they note that different logics result from different answers to this question. Thus they conclude:

As with unsettledness in general . . . , so too with the current topic: the question is ultimately one of **utility** (Beall and Restall: 29).

¹²² Beall and Restall, too, view logical pluralism as a manifestation of freedom.

From our perspective, this conclusion is overly uncritical. The question which “cases” (or which kinds of models) are relevant to logic is primarily a theoretical, not a pragmatic question, one that calls for a philosophical investigation of foundational issues central to logic. We are talking not about a foundationalist investigation, which would be futile, but about a holistic foundational investigation. Such an investigation is not intended to settle everything about logic. But the options it leaves open, the plurality it sanctions, are different, or at least differently justified, from those sanctioned by logical pluralists. For example, to the extent that intuitionistic logic is viewed as grounding logic not in the world but exclusively in the mind (in mental constructions), we, unlike most logical pluralists, reject this version of intuitionistic logic.

Do we reject the simultaneous acceptance of conflicting logics under all circumstances? No. Some circumstances under which the acceptance of such logics is unobjectionable, or is even justified, from our perspective, are (i) *human ignorance*, (ii) *formal diversity of the world*, (iii) *unavailability of a unified theory*, and (iv) *non-genuine conflict*.

- (i) *Ignorance*. First, we recognize the possibility of ignorance concerning which formal laws govern reality. If this is the case, there is room for competing conjectures (e.g., both bivalent and trivalent) about the identity of these laws.
- (ii) *Diversity*. We also recognize the possibility that the world is formally *disunified* (*diverse, chaotic*), at least to some degree. If this possibility materializes, it has significant ramifications for logic. If, for example, the formal behavior of objects on one scale (say, the microscopic scale) turns out to be governed by different laws than those governing objects on another scale (say, the macroscopic scale), then veridicality considerations will lead us to admit two logics, each for a limited domain. Each such logic will be more general, have a greater degree of necessity, etc., than the physical, biological, and other theories in its sphere of application, and as such it will retain its strong generality and modal force relative to these theories. But it will not be universal in the way a single logic could be.
- (iii) *Unavailability of a unified theory*. Suppose the world is formally unified, and we humans believe (conjecture) that it is formally unified, but we do not know how to construct a single theory of formal structure that matches its unity. In this case, we will be better off (at least temporarily) with two partial theories of formal structure, hence two logics (e.g., classical logic and fuzzy logic, classical logic and quantum mechanical logic), applicable in different contexts, than with no logic at all.¹²³

¹²³ A related possibility would be something analogous to the duality of particle and wave in physics, as viewed by Einstein and Infeld:

It seems as though we must use sometimes the one theory [of light] and sometimes the other, while at times we may use either... We have two contradictory pictures of reality: separately neither of them fully explains the phenomena of light, but together they do! (Einstein and Infeld 1938: 278).

- (iv) *Non-genuine conflict*. Finally, some cases of non-rival logics might be mistakenly viewed as cases of rivalry. Consider the case of mathematical logic and modal logic. While some might view modal logic as competing with mathematical logic over admissible logical constants, we regard it as falling under the category of “additional logics” and as such as adding a new dimension to mathematical logic rather than conflicting with it. Similarly, while relevance logic is often viewed as conflicting with mathematical logic, from our perspective it sets additional, extraformal constraints on “correct inference”, beyond the formal constraints set by mathematical logic, joining forces with it to create a more restrictive theory of inference. None of these cases is one in which we simultaneously endorse genuinely rival logics, i.e., none involves the endorsement of two theories that differ on the logical status of a purported consequence *in the same context and from the same perspective*.

The values cherished by most logical pluralists—charity, generosity, “let a thousand flowers bloom”—are humanistic values that have their correlates in the general epistemic sphere: open-mindedness, willingness to experiment, broad perspective, non-dogmatic evaluation of both existent and new theories, and so on. But truth, critical justification, and taking epistemic responsibility for our theories are also deeply humanistic values. Logic, like any other field of knowledge, is subject to powerful veridicality constraints, and one should always be careful not to confuse tolerance toward new ideas and new perspectives with tolerance toward error. Pluralism of type (b) can easily slide into radical pluralism, a pluralism that neglects the values of truth and critical examination (pluralism of type (a)). The specter of such pluralism is *frictionless theorizing*, and this we emphatically reject.

Conclusion

Toward Freedom

Knowledge requires both freedom and friction: freedom to set epistemic goals and create epistemic norms; friction generated by the targets of these goals and the demands of those norms. Every step in knowledge, from making a simple conjecture to launching a scientific revolution, involves both freedom and friction. Making a conjecture is an act of freedom, but it sets in motion a succession of acts of freedom and friction: designing and conducting tests, receiving test results, making revisions, designing and conducting new tests, and so on. Similarly, launching an epistemic revolution (scientific, mathematical, philosophical) is an act of freedom, but it is normally motivated by past frictions, and it gives rise to new challenges and new demands, hence new frictions. The actual process of constructing a theoretical body of knowledge is highly intricate, but the point that both freedom and friction are needed for knowledge is simple. Knowledge requires both freedom and friction, and an adequate theory of knowledge must account for both.

The present essay accounts for freedom and friction in two ways. First, all the major structures delineated in this work are characterized by the *complementarity* of friction and freedom. Second, within these structures, the work focuses on the application of *central friction requirements*, such as veridicality and substantiveness. Thus, the epistemic methodology developed in Part I is both holistic and foundational, but it explains how the holistic method can be used in the service of the foundational project. The model of knowledge delineated in Part II is a center-periphery model, but unlike past center-periphery models it applies the periphery norms to *all* disciplines, and it explains how disciplines that, in the past, were thought to be limited to the center (and possibly intermediate zones) can have a solid presence in the periphery. The theory of truth developed in Part III is both substantivist and flexible, both a robust correspondence theory and one that sanctions the human construction of manifold correspondence routes from theory to reality. This enables it to apply substantive correspondence standards to *all* fields of knowledge, including fields that have rarely been viewed as subject to such standards, such as logic, and fields that posed serious problems for earlier correspondence theories, such as mathematics. Finally, the foundation for logic outlined in Part IV describes the logical method as emerging from certain human needs, yet needs whose satisfaction requires a grounding of logic in *reality*. The foundational outline centers on the questions of why logic must be

grounded in reality, how it is grounded in reality, what are some of the consequences of grounding it in reality, and what are the theoretical (explanatory, problem solving) advantages of this grounding.

Our next task is to study the role of *epistemic freedom* within the general epistemic structures delineated in the present essay. It is hard to tell in advance which aspects of epistemic freedom will turn out to be most important (fruitful, revealing, explanatory). One promising topic of study is intellect's contribution to knowledge, including discovery, when this study is unencumbered by the customary assumptions that intellect's main contribution to knowledge is purely apriori, or purely pragmatic, or reducible to sensory input-output.

In setting out to investigate epistemic freedom it might be well to focus on those cases that, in philosophy, are commonly associated with its complement, epistemic friction. Just as in the present essay we have focused on cases that are commonly exempt from friction requirements (grounding in reality, substantiveness), so we might study epistemic freedom by focusing on cases that are ordinarily not associated with it. We might focus on the role of epistemic freedom in ontology and realism, we might rethink indeterminacy from the perspective of cognitive-linguistic freedom. We might investigate the role of mind in logic *in light of* its grounding in reality. We might take advantage of the new options created by our new perspective to tackle recalcitrant philosophical problems (for example, the problem of truth through scientific change or pessimistic meta-induction), to explore uncharted areas of philosophical investigation and ones that are still not adequately understood (for example, truth in ethics), and so on.

One trap that philosophers of epistemic freedom are in danger of falling into is *idealism*. This trap awaits any epistemologist who emphasizes the complementarity of mind and world, and philosophers from Kant to McDowell have been accused of falling into it. Our decision to focus on epistemic friction prior to epistemic freedom was partly motivated by the desire to avoid this trap. By approaching epistemic freedom from a point of view that emphasizes the centrality of friction, we are unlikely either to exaggerate the role played by the mind in knowledge or to neglect the role played by the world. Epistemic freedom, from our perspective, is first and foremost freedom to seek new cognitive routes to the world, test our theories against the world, develop veridical and other substantivist norms, adjust them in light of various frictions, and so on. By having the friction requirements in view at all times, we significantly diminish the danger of overlooking the world.

Another challenge facing us is the classical challenge of explaining not just how mind and world, each by itself, partakes in the project of knowledge, but how they can and do cooperate in generating knowledge. We have taken several steps in this direction in the present work (both in delineating our general epistemic structures and in working out some of their details), but much work is left to be done. Quite a few philosophers, however, are skeptical of philosophy's ability to meet this challenge, and this skepticism has led some to excessive pragmatism, others to quietism.

In my view these attitudes are unwarranted. In freeing ourselves from the traditional paradigms of apriority, analyticity, and Platonism on the one hand, radical empiricism, conventionalism, and nominalism on the other; by moving beyond thing-in-itself, God's eye view, foundationalism, coherentism, deflationism, naive correspondence, a static model of knowledge, and a total ban on circularity; by devising new tools such as foundational holism, a dynamic center-periphery model, manifold correspondence, and formal laws as a ground for logic, we create resources for meeting this challenge on a new footing. A string of three words from Wittgenstein, "Look and see", is our guide. Look and see how mind and world work in tandem. Make use of freedom. Figure out. Investigate.

References

- Allison, H. E. 1973. *The Kant-Eberhard Controversy*. Baltimore: Johns Hopkins.
- Almog, J. 1989. "Logic and the World". *Journal of Philosophical Logic* 18: 197–220.
- Alston, W. 1976. "Has Foundationalism Been Refuted?", in *Epistemic Justification*. Ithaca: Cornell University Press (1989), pp. 39–56.
- Alston, W. 1979. "Yes, Virginia, There Is a Real World". *Proceedings of the APA* 52: 779–808.
- Alston, W. 1986. "Epistemic Circularity". *Philosophy and Phenomenological Research* 47: 1–30.
- Anderson, P. W. 1972. "More Is Different". *Science* 177: 393–6.
- Aquinas, T. 1265–9. *Summa Theologica I*. New York: Benziger (1947).
- Aristotle. 4th century BC. *Metaphysics*, in *The Basic Works of Aristotle*, R. McKeon (ed.). New York: Random House (1941).
- Austin, J. L. 1950. "Truth". *Proceedings of the Aristotelian Society*, Supplement Vol. 24: 111–28.
- Azzouni, J. 2004. *Deflating Existential Consequence*. Oxford: Oxford University Press.
- Baghramian, M. 2008. "'From Realism Back to Realism': Putnam's Long Journey". *Philosophical Topics* 36: 17–35.
- Baker, G. P. and Hacker, P. M. S. 1980. *Wittgenstein*. Oxford: Blackwell.
- Balaguer, M. 1998. *Platonism and Anti-Platonism in Mathematics*. Oxford: Oxford University Press.
- Balaguer, M. 2004/9. "Platonism in Metaphysics", in *Stanford Encyclopedia of Philosophy* (Spring 2014 ed.), Edward N. Zalta (ed.). <<http://plato.stanford.edu/archives/spr2014/entries/platonism/>>.
- Bar-Elli, G. 2010. "Analyticity and Justification in Frege". *Erkenntnis* 73: 165–84.
- Barnard, R. and Horgan, T. 2006. "Truth as Mediated Correspondence". *Monist* 89: 31–50.
- Barnard, R. and Horgan, T. 2013. "The Synthetic Unity of Truth", in *Truth and Pluralism*, N. J. L. I. Pedersen and C. D. Wright (eds). Oxford: Oxford University Press, pp. 180–96.
- Barrett, R. B. and Gibson, R. F. (eds). 1990. *Perspectives on Quine*. Cambridge: Basil Blackwell.
- Barwise, J. 1972. "Absolute Logics and $L_{\infty\omega}$ ". *Annals of Pure and Applied Logic* 4: 309–40.
- Barwise, J. 1985. "Model-Theoretic Logics: Background and Aims", in *Model-Theoretic Logics*, J. Barwise and S. Feferman (eds). New York: Springer-Verlag, pp. 3–23.
- Barwise, J. and Feferman, S. (eds). 1985. *Model-Theoretic Logics*. New York: Springer-Verlag.
- Bealer, G. 1987. "The Limits of Scientific Essentialism". *Philosophical Perspectives* 1: 289–365.
- Bealer, G. 1996. "A Priori Knowledge and the Scope of Philosophy". *Philosophical Studies* 81: 121–42.
- Beall, J. C. and Restall, G. 2006. *Logical Pluralism*. Oxford: Oxford University Press.
- Beck, A. 2011. *A General Theory of Formality*. PhD thesis, University of California, San Diego.
- Belnap, N. D. 1962. "Tonk, Plonk and Plink". *Analysis* 22: 130–4.
- Ben-Menahem, Y. 1995. "Pragmatism and Revisionism: James's Conception of Truth". *Philosophical Studies* 3: 270–89.
- Benacerraf, P. 1965. "What Numbers Could Not Be". *Philosophical Review* 74: 47–73.
- Benacerraf, P. 1973. "Mathematical Truth". *Journal of Philosophy* 70: 661–79.

- Blackburn, S. 1984. *Spreading the Word*. Oxford: Oxford University Press.
- Boghossian, P. A. 1996. "Analyticity Reconsidered". *Noûs* 30: 360–91.
- Boghossian, P. A. 2000. "Knowledge of Logic", in *New Essays on the A Priori*, P. Boghossian and C. Peacocke (eds). Oxford: Oxford University Press, pp. 229–54.
- Boghossian, P. and Peacocke, C. (eds). 2000. *New Essays on the A Priori*. Oxford: Oxford University Press.
- BonJour, L. 1985. *The Structure of Empirical Knowledge*. Cambridge, Mass.: Harvard University Press.
- BonJour, L. 1994. "Against Naturalized Epistemology". *Midwest Studies in Philosophy* 19: 283–300.
- BonJour, L. 1998. *In Defense of Pure Reason*. Cambridge: Cambridge University Press.
- Bonnay, D. 2008. "Logicality and Invariance". *Bulletin of Symbolic Logic* 14: 29–68.
- Boolos, G. 1985. "Nominalist Platonism". *Philosophical Review* 94: 327–44.
- Boyd, R. 1989. "What Realism Implies and What it Does Not". *Dialectica* 43: 5–29.
- Boyd, R. 2002. "Scientific Realism", in *Stanford Encyclopedia of Philosophy*. (Spring 2010 ed.), Edward N. Zalta (ed.). <<http://plato.stanford.edu/archives/sum2010/entries/scientific-realism/>>.
- Brandom, R. B. 1994a. *Making It Explicit*. Cambridge, Mass.: Harvard University Press.
- Brandom, R. B. 1994b. "Review of *Renewing Philosophy*". *Journal of Philosophy* 91: 140–3.
- Brandom, R. B. 2000. *Articulating Reasons*. Cambridge, Mass.: Harvard University Press.
- Bueno, O. and Shalkowski, S. 2013. "Logical Constants: A Modalist Approach". *Noûs* 47: 1–24.
- Burge, T. 1998. "Frege on Knowing the Foundation". *Mind* 107: 305–47.
- Burgess, J. 1977. "Forcing", in *Handbook of Mathematical Logic*, ed. J. Barwise. Amsterdam: Elsevier, pp. 403–52.
- Burgess, J. 1983. "Why I Am Not a Nominalist". *Notre Dame Journal of Formal Logic* 24: 93–105.
- Burgess, J. and Rosen, G. 1997. *A Subject with No Object*. Oxford: Oxford University Press.
- Carnap, R. 1928. *The Logical Structure of the World*. Berkeley: University of California Press (1967).
- Carnap, R. 1934. *The Logical Syntax of Language*. London: Routledge and Kegan Paul (1937).
- Carnap, R. 1939. *Foundations of Logic and Mathematics*, *International Encyclopedia of Unified Science*, O. Neurath et al. (eds), vol. I, no. 3. Chicago: University of Chicago Press.
- Carnap, R. 1950. "Empiricism, Semantics, and Ontology". *Meaning and Necessity*. Chicago: Chicago University Press (1956), pp. 205–21.
- Carnap, R. 1963. "Replies and Systematic Expositions", in *The Philosophy of Rudolf Carnap*, P. A. Schilpp (ed.). La Salle: Open Court, pp. 859–1013.
- Cartwright, N. 1983. *How the Laws of Physics Lie*. Oxford: Oxford University Press.
- Cartwright, N. 1999. *The Dappled World*. Cambridge: Cambridge University Press.
- Casullo, A. 2015. "Four Challenges to the A Priori—A Posteriori Distinction". *Synthese* 192: 2701–24.
- Cavallès, J. 1946/60. "On Logic and the Theory of Science", in *Phenomenology and the Natural Sciences*, J. J. Kockelmans and T. J. Kiesel (eds). Evanston: Northwestern University Press (1970), pp. 357–409.
- Chateaubriand, O. 2001. *Logical Forms*. Part I. Campinas, Brazil: UNICAMP (Centro de Lógica, Epistemologia e História da Ciência).
- Chateaubriand, O. 2003. "Quine and Ontology". *Principia* 7: 41–74.

- Chateaubriand, O. 2013. "Logical Truths and Logical Facts", in *Truth*, M. Dimitru and G. Sandu (eds). Bucharest: Bucharest University Press, pp. 101–11.
- Chomsky, N. 1968. "Quine's Empirical Assumptions". *Synthese* 19: 53–68.
- Chomsky, N. 1988. "Language and Interpretation: Philosophical Reflections and Empirical Inquiry", in *Inference, Explanation, and Other Frustrations*, J. Earman (ed.). Berkeley: University of California Press (1992), pp. 99–128.
- Cohen, S. 2002. "Basic Knowledge and the Problem of Easy Knowledge". *Philosophy and Phenomenological Research* 65: 309–29.
- Comte, A. 1830. *Cours de Philosophie Positive*, vol. 1. Paris: Schleichter (1907).
- Daniels, N. 1979. "Wide Reflective Equilibrium and Theory Acceptance in Ethics". *Journal of Philosophy* 76: 256–82.
- Darwin, C. 1859. *The Origin of Species by Natural Selection*. New York: Modern Library (1993).
- David, M. 1994. *Correspondence and Disquotatation*. Oxford: Oxford University Press.
- Davidson, D. 1969. "True to the Facts". *Journal of Philosophy* 64: 691–703.
- Davidson, D. 1980. *Essay on Actions and Events*. Oxford: Oxford University Press.
- Davidson, D. 1981. "A Coherence Theory of Truth and Knowledge", in Davidson, *Subjective, Intersubjective, Objective*. Oxford: Oxford University Press (2001), pp. 137–53.
- Davidson, D. 1988. "Epistemology and Truth", in Davidson (2001), pp. 177–91.
- Davidson, D. 1990. "The Structure and Content of Truth". *Journal of Philosophy* 87: 279–328.
- Davidson, D. 1996. "The Folly of Trying to Define Truth". *Journal of Philosophy* 93: 263–78.
- Davidson, D. 2001. *Subjective, Intersubjective, Objective*. Oxford: Oxford University Press.
- Descartes, R. 1637. *Discourse on Method*. Indianapolis: Hackett (1998).
- Descartes, R. 1641. *Meditations on First Philosophy*. Indianapolis: Hackett (1998).
- Devitt, M. 1984/91. *Realism and Truth*. Oxford: Blackwell.
- Dewey, J. 1938. *Logic*. New York: Holt.
- Dilman, I. 1984. *Quine on Ontology, Necessity, and Experience*. Albany: State University of New York.
- Dreben, B. 1992. "Putnam, Quine—and the Facts". *Philosophical Topics* 20: 293–315.
- Duhem, P. 1906/14. *The Aim and Structure of Physical Theory*. Princeton: Princeton University Press (1954).
- Dummett, M. 1963. "Realism", in Dummett, *Truth and Other Enigmas*. Cambridge, Mass.: Harvard University Press (1978), pp. 146–65.
- Dummett, M. 1973. "The Significance of Quine's Indeterminacy Thesis", in Dummett (1978): 375–419.
- Dummett, M. 1973/81. *Frege*. New York: Harper and Row.
- Dummett, M. 1976/91. *The Logical Basis of Metaphysics*. Cambridge, Mass.: Harvard University Press.
- Dummett, M. 1978. *Truth and Other Enigmas*. Cambridge, Mass.: Harvard University Press.
- Dupré, J. 1993. *The Disorder of Things*. Cambridge, Mass.: Harvard University Press.
- Dutilh Novaes, C. 2011. "The Different Ways in which Logic is (said to be) Formal". *History and Philosophy of Logic* 32: 303–32.
- Dutilh Novaes, C. 2012. "Reassessing Logical Hylomorphism and the Demarcation of Logical Constants". *Synthese* 185: 387–410.
- Dworkin, R. 1996. "Objectivity and Truth: You'd Better Believe it". *Philosophy and Public Affairs* 25: 87–139.

- Dyson, F. 1988. *Infinite in All Directions*. New York: Harper and Row.
- Eco, U. 1995. "A View from Elm Street" in *On Quine*, P. Leonardi and M. Santambrogio (eds). Cambridge: Cambridge University Press, pp. 22–36.
- Einstein, A. and Infeld, L. 1938. *The Evolution of Physics*. Cambridge: Cambridge University Press.
- Elgin, C. Z. 1996. *Considered Judgment*. Princeton: Princeton University Press.
- Enderton, H. B. 1972/2001. *A Mathematical Introduction to Logic*. San Diego: Harcourt.
- Esfeld, M. 2001. "Gonseth and Quine". *Dialectica* 55: 199–219.
- Etchemendy, J. 1990. *The Concept of Logical Consequence*. Cambridge, Mass.: Harvard University Press.
- Feferman, S. 1984. "Foundational Ways", in Feferman, *In Light of Logic*, Oxford: Oxford University Press (1998), pp. 94–104.
- Feferman, S. 1993a. "Working Foundations—'91", in Feferman (1998), pp. 105–24.
- Feferman, S. 1993b. "Why a Little Bit Goes a Long Way: Logical Foundations of Scientifically Applicable Mathematics", in Feferman (1998), pp. 284–98.
- Feferman, S. 1998. *In Light of Logic*. Oxford: Oxford University Press.
- Feferman, S. 1999. "Logic, Logics, and Logicism". *Notre Dame Journal of Formal Logic* 40: 31–54.
- Feferman, S. 2000. "Mathematical Intuition vs. Mathematical Monsters". *Synthese* 125: 317–22.
- Feferman, S. 2010. "Set-Theoretical Invariance Criteria for Logicality". *Notre Dame Journal of Formal Logic* 51: 3–20.
- Feferman, S. 2015. "Which Quantifiers are Logical? A Combined Semantical and Inferential Criterion", in *Quantifiers, Quantifiers, and Quantifiers: Themes in Logic, Metaphysics, and Language*, A. Torza (ed.). Switzerland: Springer, pp. 19–30.
- Feferman, S. and Hellman, G. 2000. "Challenges to Predicative Foundations of Arithmetic", in *Between Logic and Intuition*, G. Sher and R. Tieszen (eds). Cambridge: Cambridge University Press, pp. 317–38.
- Feynman, R. 1965. *The Character of Physical Law*. Cambridge, Mass.: MIT Press (1967).
- Field, H. 1972. "Tarski's Theory of Truth". *Journal of Philosophy* 69: 347–75.
- Field, H. 1980. *Science without Numbers*. Oxford: Basil Blackwell.
- Field, H. 1986. "The Deflationary Conception of Truth", in *Fact, Science and Morality*, G. MacDonald and C. Wright (eds). Oxford: Blackwell, pp. 55–117.
- Field, H. 1989. *Realism, Mathematics and Modality*. Oxford: Basil Blackwell.
- Field, H. 1994. "Deflationist Views of Meaning and Content". *Mind* 103: 249–85.
- Field, H. 2008. *Saving Truth from Paradox*. Oxford: Oxford University Press.
- Field, H. 2009a. "What is the Normative Role of Logic?". *Proceedings of the Aristotelian Society*, Supplementary Vol. 83: 251–68.
- Field, H. 2009b. "Pluralism in Logic". *Review of Symbolic Logic* 2: 342–59.
- Fine, K. 1995. "Ontological Dependence". *Proceedings of the Aristotelian Society* 95: 269–90.
- Fine, K. 2001. "The Question of Realism". *Philosophers' Imprint* 1: 1–30.
- Fine, K. 2005. "Our Knowledge of Mathematical Objects", in *Oxford Studies in Epistemology*, vol. 1, T. Szabó Gendler and J. Hawthorne (eds). Oxford: Oxford University Press, pp. 89–109.
- Fine, K. 2012. "Guide to Ground", in *Metaphysical Grounding*, F. Correia and B. Schnieder (eds). Cambridge: Cambridge University Press, pp. 37–80.

- Finkelstein, D. 1969. "Matter, Space and Logic", in *Boston Studies in the Philosophy of Science*, vol. 5, R. S. Cohen and M. W. Wartofsky (eds). Dordrecht: D. Reidel, pp. 199–215.
- Floyd, J. 1995. "On Saying What You Really Want to Say: Wittgenstein, Gödel, and the Trisection of the Angle", in *Essays on the Development of the Foundations of Mathematics*, J. Hintikka (ed.). Dordrecht: Kluwer, pp. 373–425.
- Fodor, J. 1974. "Special Sciences". *Synthese* 28: 77–115.
- Fodor, J. and Lepore, E. 1992. *Holism*. Oxford: Basil Blackwell.
- Fraenkel, A. A. et al. 1958/73. *Foundations of Set Theory*. Amsterdam: North Holland.
- Frege, G. 1879. "Begriffsschrift", in *From Frege to Gödel*, J. van Heijenoort (ed.). Cambridge, Mass.: Harvard University Press (1967), pp. 5–82.
- Frege, G. 1884. *The Foundations of Arithmetic*. Evanston: Northwestern, 1968.
- Frege, G. 1892. "On Sense and Reference", in *Translations from the Philosophical Writings of Gottlob Frege*. Oxford: Basil Blackwell (1970), pp. 56–78.
- Frege, G. 1893. *The Basic Laws of Arithmetic I*. Berkeley: University of California Press (1964).
- Frege, G. 1918. "Thoughts", in *Logical Investigations*. Oxford: Basil Blackwell (1977), pp. 1–30.
- Frege, G. 1919. "The Thought: A Logical Inquiry". *Mind* 65 (1956): 289–311.
- Freud, S. 1940. *An Outline of Psycho-Analysis*. New York: Norton (1969).
- Friedman, M. 1974. "Explanation and Scientific Understanding", *Journal of Philosophy* 71: 5–19.
- Friedman, M. 1997. "Philosophical Naturalism". *Proceedings of the APA* 71: 7–21.
- Friedman, M. 1999. *Reconsidering Logical Positivism*. Cambridge: Cambridge University Press.
- Friedman, M. 2001. *Dynamics of Reason*. Stanford: CSLI Publications.
- Gibbard, A. 1990. *Wise Choices, Apt Feelings*. Cambridge, Mass.: Harvard University Press.
- Gibson, R. F. 1982. *The Philosophy of W.V. Quine*. Tampa: University Presses of Florida.
- Gibson, R. F. 1988. *Enlightened Empiricism*. Tampa: University Presses of Florida.
- Glymour, C. 1980. *Theory and Evidence*. Princeton: Princeton University Press.
- Gochet, P. 1986. *Ascent to Truth*. Munich: Philosophia Verlag.
- Gödel, K. 1940. "The Consistency of the Axiom of Choice and of the Generalized Continuum Hypothesis with the Axioms of Set Theory", in *Collected Works*, vol. II. Oxford: Oxford University Press (1990), pp. 33–101.
- Gödel, K. 1953–9. "Is Mathematics Syntax of Language?", in *Collected Works*, vol. III. Oxford: Oxford University Press (1995), pp. 334–62.
- Goldfarb, W. 2010. "Frege's Conception of Logic", in *The Cambridge Companion to Frege*, M. Potter and T. Ricketts (eds). Cambridge: Cambridge University Press, pp. 63–85.
- Goldman, A. 1979. "What Is Justified Belief?", in *Justification and Knowledge*, G. S. Pappas (ed.). Dordrecht: Reidel, pp. 1–23.
- Gómez-Torrente, M. 2002. "The Problem of Logical Constants". *Bulletin of Symbolic Logic* 8: 1–37.
- Gómez-Torrente, M. 2003. "The 'Must' and the 'Heptahedron'. Remarks on Remarks". *Theoria* 18: 199–206.
- Goodman, N. 1953/5. "The New Riddle of Induction", in *Fact, Fiction, and Forecast*. Indianapolis: Bobbs-Merrill, pp. 59–83.
- Goodman, N. 1976. *Languages of Art*. Indianapolis: Hackett.
- Grayling, A. C. 2001. "Understanding Realism". <<http://www.acgrayling.com/understanding-realism>>.

- Grice, H. P. and P. F. Strawson. 1956. "In Defence of a Dogma". *Philosophical Review* 65: 141–58.
- Grünbaum, A. 1960. "The Duhemian Argument", in *Can Theories be Refuted?* S. D. Harding (ed.). Dordrecht: D. Reidel (1976), pp. 116–31.
- Grünbaum, A. 1971. "Is it *Never* Possible to Falsify a Hypothesis Irrevocably?", in Harding (1976), pp. 260–88.
- Gupta, A. 1993a. "Minimalism". *Philosophical Perspectives* 7: 359–69.
- Gupta, A. 1993b. "A Critique of Deflationism". *Philosophical Topics* 21: 57–81.
- Gupta, A. and Belnap, N. 1993. *The Revision Theory of Truth*. Cambridge, Mass.: MIT Press.
- Haack, S. 1993. *Evidence and Inquiry*. Oxford: Blackwell.
- Hacking, I. 1996. "The Disunities of the Sciences", in *The Disunity of Science*, P. Galison and D. J. Stump (eds). Stanford: Stanford University Press, pp. 37–74.
- Hahn, L. E. and Schilpp, P. A. (eds). 1986. *The Philosophy of W. V. Quine*. La Salle: Open Court.
- Hale, B. and Wright, C. 2001. *The Reason's Proper Study*. Oxford: Oxford University Press.
- Hanna, R. 2000. "Kant, Truth and Human Nature". *British Journal for the History of Philosophy* 8: 225–50.
- Hanna, R. 2001. *Kant and the Foundations of Analytic Philosophy*. Oxford: Oxford University Press.
- Hanna, R. 2004/13. "Kant's Theory of Judgment", in *Stanford Encyclopedia of Philosophy* (Spring 2014 ed.), Edward N. Zalta (ed.). <<http://plato.stanford.edu/archives/sum2014/entries/kant-judgment/>>.
- Hanna, R. 2006. *Rationality and Logic*. Cambridge, Mass.: MIT Press.
- Hanson, W. H. 1997. "The Concept of Logical Consequence". *Philosophical Review* 106: 365–409.
- Hanson, W. H. 2002. "The Formal-Structural View of Logical Consequence: A Reply to Gila Sher". *Philosophical Review* 111: 243–58.
- Harding, S. G. (ed.). 1976. *Can Theories Be Refuted?* Dordrecht: D. Reidel.
- Harman, G. 1986. *Change in View*. Cambridge, Mass.: MIT Press.
- Harman, G. 2009. "Field on the Normative Role of Logic". *Proceedings of the Aristotelian Society* 109: 333–5.
- Hawking, S. W. 1980. *Is the End in Sight for Theoretical Physics?* Cambridge: Cambridge University Press.
- Hawthorne, J. 2007. "A Priority and Externalism" in *Internalism and Externalism in Semantics and Epistemology*, S. Goldberg (ed.). Oxford: Oxford University Press, pp. 201–18.
- Hellman, G. 1989. *Mathematics Without Numbers*. Oxford: Oxford University Press.
- Hempel, C. 1965. *Aspects of Scientific Explanation*. New York: Free Press.
- Hempel, C. 1966. *Philosophy of Natural Science*. Englewood Cliffs: Prentice Hall.
- Hempel, C. and Oppenheim, P. 1948. "Studies in the Logic of Explanation". *Philosophy of Science* 15: 135–75.
- Henkin, L. 1959. "Some Remarks on Infinitely Long Formulas" in *Infinitistic Methods*. Oxford: Pergamon (1961), pp. 167–83.
- Herburt, G. K. 1959. "The Analytic and the Synthetic". *Philosophy of Science* 26: 104–13.
- Hesse, M. 1970. "Duhem, Quine and a New Empiricism", in *Can Theories Be Refuted?* S. G. Harding (ed.) (1976), pp. 184–204.

- Higginbotham, J. and May, R. 1981. "Questions, Quantifiers and Crossing". *Linguistic Review* 1: 41–79.
- Hintikka, J. 1995. "What Is Elementary Logic? Independence-Friendly Logic as the True Core Arc of Logic", in *Physics, Philosophy, and the Scientific Community*, K. Gavroglu et al. (eds). Dordrecht: Kluwer, pp. 301–26.
- Hodes, H. 1984. "Logicism and the Ontological Commitments of Arithmetic". *Journal of Philosophy* 81: 123–49.
- Hodes, H. 1990. "Where Do the Natural Numbers Come From". *Synthese* 84: 347–407.
- Hofstadter, A. 1954. "The Myth of the Whole: A Consideration of Quine's View of Knowledge". *Journal of Philosophy* 51: 397–417.
- Hofweber, T. 2005. "Number Determiners, Numbers, and Arithmetic". *Philosophical Review* 114: 179–225.
- Hookway, C. 1988. *Quine*. Stanford: Stanford University Press.
- Horgan, T. 2001. "Contextual Semantics and Metaphysical Realism: Truth as Indirect Correspondence", in *The Nature of Truth*, M. Lynch (ed.). Cambridge, Mass.: MIT Press, pp. 67–95.
- Horwich, P. 1990/8. *Truth*. Oxford: Oxford University Press. (References to the 1990 edition.)
- Hughes, R. I. G. 1989. *The Structure and Interpretation of Quantum Mechanics*. Cambridge, Mass.: Harvard University Press.
- Hume, D. 1739–40. *A Treatise of Human Nature*. Oxford: Oxford University Press (2000).
- Hylton, P. 2007. *Quine*. New York: Routledge.
- James, W. 1890. "The Stream of Thought", in *The Principles of Psychology*, vol. 1. New York: Dover (1950), pp. 224–90.
- James, W. 1907. "Pragmatism", in *Pragmatism and Other Writings*. New York: Penguin (2000), pp. 1–132.
- Jech, T. J. 2003. *Set Theory*. Berlin: Springer.
- Jenkins, C. S. 2005. "Realism and Independence". *American Philosophical Quarterly* 42: 199–209.
- Jenkins, C. S. 2008. *Grounding Concepts*. Oxford: Oxford University Press.
- Kant, I. 1770s–1800. *Lectures on Logic*. Cambridge: Cambridge University Press (1992).
- Kant, I. 1781/7. *Critique of Pure Reason*, N. Kemp Smith (trans.), London: Macmillan (1929); P. Guyer and A. W. Wood (trans.), Cambridge: Cambridge University Press (1998). (Unless otherwise indicated, citations are based on the Kemp Smith translation.)
- Kant, I. 1783. *Prolegomena to Any Future Metaphysics*. New York: Macmillan (1950).
- Keenan, E. L. and Stavi, J. 1986. "A Semantic Characterization of Natural Language Determiners". *Linguistics and Philosophy* 9: 253–329.
- Keisler, H. J. 1970. "Logic with the Quantifier 'There Exist Uncountably Many'". *Annals of Mathematical Logic* 1: 1–93.
- Khlentzos, D. 2011/11. "Challenges to Metaphysical Realism" in *Stanford Encyclopedia of Philosophy*, Edward N. Zalta (ed.). <<http://plato.stanford.edu/archives/spr2011/entries/realism-sem-challeng/>>.
- Kim, J. 1988. "What is Naturalized Epistemology?" *Philosophical Perspectives* 2: 381–406.
- Kirkham, R. L. 1992. *Theories of Truth*. Cambridge, Mass.: MIT Press.
- Kitcher, P. 1981. "Explanatory Unification", *Philosophy of Science* 48: 507–31.
- Kitcher, P. 1989. "Explanatory Unification and the Causal Structure of the World", in *Scientific Explanation*, P. Kitcher and W. Salmon (eds). Minneapolis: Minnesota University Press, pp. 410–505.
- Kitcher, P. 2001. *Science, Truth, and Democracy*. Oxford: Oxford University Press.

- Korsgaard, C. 1996. *The Sources of Normativity*. Cambridge: Cambridge University Press.
- Kreisel, G. 1967. "Informal Rigor and Completeness Proofs", in *Problems in the Philosophy of Mathematics*, I. Lakatos (ed.). Amsterdam: North Holland, pp. 138–71.
- Kripke, S. 1972/80. *Naming and Necessity*. Cambridge, Mass.: Harvard University Press.
- Kripke, S. 1975. "Outline of a Theory of Truth". *Journal of Philosophy* 72: 690–716.
- Lammenranta, M. 1996. "Reliabilism and Circularity". *Philosophy and Phenomenological Research* 56: 111–24.
- Lammenranta, M. 2009. "Epistemic Circularity". *Internet Encyclopedia of Philosophy*, ISSN 2161-0002 <<http://www.iep.utm.edu/ep-circ/>>, 2015.
- Laughlin, R. B. and Pines, D. 2000. "The Theory of Everything". *Proceedings of the National Academy of Sciences of the USA* 97: 28–31.
- Laughlin, R. B. et al. 2000. "The Middle Way". *Proceedings of the National Academy of Sciences of the USA* 97: 32–7.
- Lear, J. 1982. "Aristotle's Philosophy of Mathematics". *Philosophical Review* 91: 161–92.
- Leonardi, P. and M. Santambrogio (eds). 1995. *On Quine.*, Cambridge: Cambridge University Press.
- Lévy, A. 1960. "Axiom Schemata of Strong Infinity in Axiomatic Set Theory". *Pacific Journal of Mathematics* 10: 23–38.
- Lindström, P. 1966. "First Order Predicate Logic with Generalized Quantifiers". *Theoria* 32: 186–95.
- Lindström, P. 1969. "On Extensions of Elementary Logic". *Theoria* 35: 1–11.
- Lindström, P. 1974. "On Characterizing Elementary Logic", in *Logical Theory and Semantic Analysis*, S. Stenlund (ed.). Dordrecht: D. Reidel, pp. 129–46.
- Linnebo, Ø. 2009/13. "Platonism in the Philosophy of Mathematics", in *Stanford Encyclopedia of Philosophy* (Winter 2013 ed.), Edward N. Zalta (ed.). <<http://plato.stanford.edu/archives/win2013/entries/platonism-mathematics/>>.
- Longueness, B. 1993/8. *Kant and the Capacity to Judge*. Princeton: Princeton University Press (1998).
- Lynch, M. 1998. *Truth in Context*. Cambridge, Mass.: MIT Press.
- Lynch, M. 2001. "A Functionalism Theory of Truth", in *The Nature of Truth*. Cambridge, Mass.: MIT Press, pp. 723–49.
- Lynch, M. 2004a. *True to Life*. Cambridge, Mass.: MIT Press.
- Lynch, M. 2004b. "Truth and Multiple Realizability". *Australasian Journal of Philosophy* 82: 384–408.
- Lynch, M. 2005. "Replies to Critics". *Philosophical Books* 46: 311–42.
- Lynch, M. 2009. *Truth as One and Many*. Oxford: Oxford University Press.
- Maaß, J. G. 1789. "On the Highest Principle of Synthetic Judgments in Relation to the Theory of Mathematical Certainty". *Philosophisches Magazin* 2 Bd. 1 St., pp. 186–231.
- Mac Lane, S. 1971. *Categories for the Working Mathematician*. New York: Springer-Verlag.
- MacFarlane, J. 2000. *What Does it Mean to Say that Logic is Formal?* PhD thesis, University of Pittsburgh.
- MacFarlane, J. 2002. "Frege, Kant, and the Logic in Logicism". *Philosophical Review* 111: 25–65.
- MacFarlane, J. 2004. "McDowell's Kantianism". *Theoria* 70: 250–65.
- MacFarlane, J. 2005. "Making Sense of Relative Truth". *Proceedings of the Aristotelian Society* 105: 321–39.

- MacFarlane, J. 2005/9. "Logical Constants", in *Stanford Encyclopedia of Philosophy*, (Fall 2015 ed.), Edward N. Zalta (ed.). <<http://plato.stanford.edu/archives/fall2015/entries/logical-constants/>>.
- Maddy, P. 1997. *Naturalism in Mathematics*. Oxford: Oxford University Press.
- Maddy, P. 2007. *Second Philosophy*. Oxford: Oxford University Press.
- Maddy, P. 2012. "The Philosophy of Logic". *Bulletin of Symbolic Logic* 18: 481–504.
- Marcus, R. B. 1988. "A Backward Look at Quine's Animadversions on Modalities", in *Perspectives on Quine*, R. B. Barrett and R. F. Gibson (eds). Cambridge: Basil Blackwell (1990), pp. 230–43.
- McCarthy, T. 1981. "The Idea of a Logical Constant". *Journal of Philosophy* 78: 499–523.
- McDermid, D. 1998. "Pragmatism and Truth: The Comparison Objection to Correspondence". *Review of Metaphysics* 51: 775–811.
- McDowell, J. 1994. *Mind and World*. Cambridge, Mass.: Harvard University Press.
- McGee, V. 1991. *Truth, Vagueness and Paradox*. Indianapolis: Hackett.
- McGee, V. 1992. "Two Problems with Tarski's Theory of Consequence". *Proceedings of the Aristotelian Society* 92: 273–92.
- McGee, V. 1996. "Logical Operations". *Journal of Philosophical Logic* 25: 567–80.
- McGee, V. 2004. "Tarski's Staggering Existential Assumptions". *Synthese* 142: 371–87.
- McGinn, C. 1983. "Review of Quine's *Theories and Things*". *Journal of Philosophy* 80: 239–46.
- Miller, A. 2002/14. "Realism", in *Stanford Encyclopedia of Philosophy*, (Winter 2014 ed.), Edward N. Zalta (ed.). <<http://plato.stanford.edu/archives/win2014/entries/realism/>>.
- Moltmann, F. 2013. "Reference to Numbers in Natural Language". *Philosophical Studies* 162: 499–536.
- Moore, G. H. 1982. *Zermelo's Axiom of Choice*. New York: Springer-Verlag.
- Morrison, M. 2000. *Unifying Scientific Theories*. Cambridge: Cambridge University Press.
- Mostowski, A. 1957. "On a Generalization of Quantifiers". *Fundamenta Mathematicae* 44: 12–36.
- Nagel, T. 1997. *The Last Word*. Oxford: Oxford University Press.
- Neurath, O. 1932. "Protocol Statements". *Philosophical Papers 1913–1946*. Dordrecht: D. Reidel (1983), pp. 91–9.
- Nicolas, D. 2013. "The Logic of Mass Expressions", in *Stanford Encyclopedia of Philosophy*, (Spring 2014 ed.), Edward N. Zalta (ed.). <<http://plato.stanford.edu/archives/spr2014/entries/logic-massexpress/>>.
- Nozick, R. 2001. *Invariances*. Cambridge, Mass.: Harvard University Press.
- Pagin, P. 1997. "Is Compositionality Compatible with Holism?". *Mind and Language* 12: 11–33.
- Parsons, C. 1979. "Mathematical Intuition". *Proceedings of the Aristotelian Society* 80: 145–68.
- Parsons, C. 1990. "The Structuralist View of Mathematical Objects". *Synthese* 84: 303–47.
- Parsons, C. 1995. "Platonism and Mathematical Intuition in Kurt Gödel's Thought". *Bulletin of Symbolic Logic* 1: 44–74.
- Parsons, C. 2000. "Reason and Intuition". *Synthese* 125: 299–315.
- Parsons, C. 2008. *Mathematical Thought and Its Objects*. Cambridge: Cambridge University Press.
- Patterson, D. 2012. *Alfred Tarski*. London: Palgrave Macmillan.
- Peacocke, C. 1976. "What Is a Logical Constant?". *Journal of Philosophy* 73: 221–40.
- Peacocke, C. 2000. "Explaining the A Priori: The Programme of Moderate Rationalism", in *New Essays on the A Priori*, P. Boghossian and C. Peacocke (eds). Oxford: Oxford University Press, pp. 255–85.
- Pedersen, N. J. L. I. 2010. "Stabilizing Alethic Pluralism". *Philosophical Quarterly* 60: 92–108.

- Pedersen, N. J. L. I. 2012a. "Recent Work on Alethic Pluralism". *Analysis* 72: 588–607.
- Pedersen, N. J. L. I. 2012b. "True Alethic Functionalism?" *Philosophical Studies* 20: 125–33.
- Pedersen, N. J. L. I. and Wright, C. D. 2012/13. "Pluralist Theories of Truth", in *Stanford Encyclopedia of Philosophy*. (Spring 2013 ed.), Edward N. Zalta (ed.). <<http://plato.stanford.edu/archives/spr2013/entries/truth-pluralist/>>.
- Pedersen, N. J. L. I. and Wright, C. D. 2013a. "Pluralism about Truth as Alethic Disjunctivism", in *Truth and Pluralism*, N. J. L. I. Pedersen and C. D. Wright (eds). Oxford: Oxford University Press (2013b), pp. 87–112.
- Pedersen, N. J. L. I. and Wright, C. D. (eds). 2013b. *Truth and Pluralism*. Oxford: Oxford University Press.
- Peters, S. and Westerståhl, D. 2006. *Quantifiers in Language and Logic*. Oxford: Oxford University Press.
- Popper, K. R. 1959. *The Logic of Scientific Discovery*. London: Routledge.
- Posy, C. 2000. "Immediacy and the Birth of Reference in Kant: The Case for Space", in *Between Logic and Intuition*, G. Sher and R. Tieszen (eds). Cambridge: Cambridge University Press, pp. 155–85.
- Prauss, G. 1969. "The Problem of Truth in Kant". *Contemporary German Philosophy*, vol. 1, D. E. Christensen (ed.). University Park: Penn State University Press (1982), pp. 91–109.
- Priest, G. 2001. "Logic: One or Many?", in *Logical Consequences*, B. Brown and J. Woods (eds). Oslo: Hermes, pp. 23–38.
- Prior, A. N. 1960. "The Runabout Inference-Ticket". *Analysis* 21: 38–9.
- Putnam, H. 1962. "The Analytic and the Synthetic", in *Mind, Language and Reality*. Cambridge: Cambridge University Press, pp. 33–69.
- Putnam, H. 1965. "Craig's Theorem", in *Mathematics, Matter and Method*. Cambridge: Cambridge University Press (1975/9), pp. 228–36.
- Putnam, H. 1967a. "The Thesis that Mathematics is Logic", in Putnam (1975/9), pp. 12–42.
- Putnam, H. 1967b. "Mathematics without Foundation", in Putnam (1975/9), pp. 43–59.
- Putnam, H. 1968. "The Logic of Quantum Mechanics", in Putnam (1975/9), pp. 174–97.
- Putnam, H. 1975/9. *Mathematics, Matter and Method*. Cambridge: Cambridge University Press.
- Putnam, H. 1981. *Reason, Truth and History*. Cambridge: Cambridge University Press.
- Putnam, H. 1983a. *Realism and Reason*. Cambridge: Cambridge University Press.
- Putnam, H. 1983b. "Convention: A Theme in Philosophy", in Putnam (1983a), pp. 170–83.
- Putnam, H. 1983c. "Why Reason Can't Be Naturalized", in Putnam (1983a), pp. 229–47.
- Putnam, H. 1987. *The Many Faces of Realism*. La Salle: Open Court.
- Putnam, H. 1994a. *Words and Life*. Cambridge, Mass.: Harvard University Press.
- Putnam, H. 1994b. "Rethinking Mathematical Necessity", in Putnam (1994a), pp. 245–63.
- Putnam, H. 1994c. "The Question of Realism", in Putnam (1994a), pp. 295–312.
- Putnam, H. 1999. *The Threefold Cord*. New York: Columbia University Press.
- Quine, W. V. 1935. "Truth by Convention", in *The Ways of Paradox and Other Essays*. Cambridge, Mass.: Harvard University Press (1976), pp. 77–106.
- Quine, W. V. 1940/51. *Mathematical Logic*. Cambridge, Mass.: Harvard University Press.
- Quine, W. V. 1948. "On What There Is", in *From a Logical Point of View*. Cambridge, Mass.: Harvard University Press (1980), pp. 1–19.
- Quine, W. V. 1950. "Identity, Ostension, and Hypostasis", in Quine (1980), pp. 65–79.

- Quine, W. V. 1951a. "Two Dogmas of Empiricism", in Quine (1980), pp. 20–46.
- Quine, W. V. 1951b. "On Carnap's Views on Ontology", in Quine (1976), pp. 203–11.
- Quine, W. V. 1952. "On Mental Entities", in Quine (1976), pp. 221–7.
- Quine, W. V. 1953a. "Logic and the Reification of Universals", in Quine (1980), pp. 102–29.
- Quine, W. V. 1953b. "Mr. Strawson on Logical Theory", in Quine (1976), pp. 137–57.
- Quine, W. V. 1953c. "Three Grades of Modal Involvement", in Quine (1976), pp. 158–76.
- Quine, W. V. 1954a. "Carnap and Logical Truth", in Quine (1976), pp. 107–32.
- Quine, W. V. 1954b. "The Scope and Language of Science", in Quine (1976), pp. 228–45.
- Quine, W. V. 1955. "Posits and Reality", in Quine (1976), pp. 246–54.
- Quine, W. V. 1960. *Word and Object*. Cambridge, Mass.: MIT Press.
- Quine, W. V. 1962. Letter to Grünbaum, in *Can Theories Be Refuted?* S. G. Harding (ed.). Dordrecht: D. Reidel (1976), p. 132.
- Quine, W. V. 1969a. "Epistemology Naturalized", in *Ontological Relativity and Other Essays*. New York: Columbia University Press (1969d), pp. 69–90.
- Quine, W. V. 1969b. "Ontological Relativity", in Quine (1969d), pp. 26–68.
- Quine, W. V. 1969c. "Existence and Quantification", in Quine (1969d), pp. 91–113.
- Quine, W. V. 1969d. *Ontological Relativity and Other Essays*. New York: Columbia University Press.
- Quine, W. V. 1970/86. *Philosophy of Logic*. Cambridge, Mass.: Harvard University Press.
- Quine, W. V. 1973. *The Roots of Reference*. La Salle: Open Court.
- Quine, W. V. 1975. "On Empirically Equivalent Systems of the World". *Erkenntnis* 9: 313–28.
- Quine, W. V. 1976. *The Ways of Paradox and Other Essays*. Cambridge, Mass.: Harvard University Press.
- Quine, W. V. 1980. *From a Logical Point of View*. Cambridge, Mass.: Harvard University Press.
- Quine, W. V. 1981a. *Theories and Things*. Cambridge, Mass.: Harvard University Press.
- Quine, W. V. 1981b. "Things and Their Place in Theories", in Quine (1981a), pp. 1–23.
- Quine, W. V. 1981c. "Responding to M. J. Cresswell", in Quine (1981a), pp. 179–81.
- Quine, W. V. 1981d. "Responding to David Armstrong", in Quine (1981a), pp. 182–4.
- Quine, W. V. 1986a. "Reply to Harold N. Lee", in *The Philosophy of W. V. Quine*, L. E. Hahn and P. A. Schillp (eds). La Salle: Open Court, pp. 315–8.
- Quine, W. V. 1986b. "Reply to Charles Parsons", in Hahn and Schillp, pp. 396–403.
- Quine, W. V. 1990. *Pursuit of Truth*. Cambridge, Mass.: Harvard University Press.
- Quine, W. V. 1991. "Two Dogmas in Retrospect". *Canadian Journal of Philosophy* 21: 265–74.
- Quine, W. V. 1995. "Reactions", in *On Quine*, P. Leonardi and M. Santambrogio (eds). Cambridge: Cambridge University Press, pp. 347–61.
- Railton, P. 2000. "A Priori Rules: Wittgenstein on the Normativity of Logic", in *New Essays on the A Priori*, P. Boghossian and C. Peacocke (eds). Oxford: Oxford University Press, pp. 170–96.
- Rawls, J. 1971. *A Theory of Justice*. Cambridge, Mass.: Harvard University Press.
- Rawls, J. 1974–5. "The Independence of Moral Theory". *Proceedings of the APA* 48: 5–22.
- Reid, T. 1785. *The Works of Thomas Reid*, vol. III. Charlestown: Samuel Etherbridge (1815).
- Rescher, N. 1973. *The Coherence Theory of Truth*. Oxford: Oxford University Press.
- Rescher, N. 2010. *Infinite Regress*. New Brunswick: Transaction.

- Resnik, M. D. 1981. "Mathematics as a Science of Patterns: Ontology and Reference". *Noûs* 15: 529–50.
- Resnik, M. D. 1982. "Mathematics as a Science of Patterns: Epistemology". *Noûs* 16: 95–105.
- Resnik, M. D. 1990. "Immanent Truth". *Mind* 99: 405–24.
- Resnik, M. D. 1997. *Mathematics as a Science of Patterns*. Oxford: Oxford University Press.
- Robinson, A. 1966/96. *Non-Standard Analysis*. Princeton: Princeton University Press.
- Romanos, G. D. 1983. *Quine and Analytic Philosophy*. Cambridge, Mass.: MIT Press.
- Rorty, R. 1979. *Philosophy and the Mirror of Nature*. Princeton: Princeton University Press.
- Rorty, R. 1989. *Contingency, Irony, and Solidarity*. Cambridge: Cambridge University Press.
- Rorty, R. 1991. "Introduction: Antirepresentationalism, Ethnocentrism, and Liberalism", in *Objectivity, Relativism, and Truth*. Cambridge: Cambridge University Press, pp. 1–17.
- Rosen, G. 1990. "Modal Fictionalism". *Mind* 99: 327–54.
- Rosen, G. 1994. "Objectivity and Modern Idealism: What is the Question?", in *Philosophy of Mind*, M. Michael and J. O'Leary-Hawthorne (eds). Dordrecht: Kluwer, pp. 277–319.
- Rosen, G. 2010. "Metaphysical Dependence: Grounding and Reduction", in *Modality, Metaphysics, Logic, and Epistemology*, B. Hale and A. Hoffmann (eds). Oxford: Oxford University Press, pp. 109–36.
- Rosenkoetter, T. 2009. "Truth Criteria and the Very Project of a Transcendental Logic". *Archiv für Geschichte der Philosophie* 91: 193–236.
- Russell, B. 1903/38. *The Principles of Mathematics*. New York: Norton.
- Russell, B. 1919. *Introduction to Mathematical Philosophy*. New York: Simon and Schuster (1971).
- Russell, B. 1938. "Introduction to the Second Edition", in *The Principles of Mathematics*. New York: Norton, pp. v–xiv.
- Russell, G. 2008. *Truth in Virtue of Meaning*. Oxford: Oxford University Press.
- Sagi, G. 2014. "Formality in Logic: From Logical Terms to Semantic Constraints". *Logique et Analyse* 57: 259–76.
- Schaffer, J. 2003. "Is There a Fundamental Level?". *Noûs* 37: 498–517.
- Schaffer, J. 2010. "The Priority of the Whole". *Philosophical Review* 119: 31–76.
- Schlick, M. 1935. "Facts and Propositions". *Analysis* 2: 65–70.
- Sellars, W. 1956. *Empiricism and the Philosophy of Mind*. Cambridge, Mass.: Harvard University Press (1997).
- Sen, A. 2002. *Rationality and Freedom*. Cambridge, Mass.: Harvard University Press.
- Shahan, R. W. and Swoyer, C. (eds). 1979. *Essays on the Philosophy of W. V. Quine*. Norman: Oklahoma University Press.
- Shapiro, S. 1991. *Foundations without Foundationalism*. Oxford: Oxford University Press.
- Shapiro, S. 1997. *Philosophy of Mathematics*. Oxford: Oxford University Press.
- Shapiro, S. 2000. "The Status of Logic", in *New Essays on the A Priori*, P. Boghossian and C. Peacocke (eds). Oxford: Oxford University Press, pp. 333–66.
- Sheffer, H. M. 1926. "Review of *Principia Mathematica*". *Isis* 8: 226–31.
- Sher, G. 1989. "A Conception of Tarskian Logic". *Pacific Philosophical Quarterly* 70: 341–68.
- Sher, G. 1991. *The Bounds of Logic*. Cambridge, Mass.: MIT Press.
- Sher, G. 1996. "Did Tarski Commit 'Tarski's Fallacy'?" *Journal of Symbolic Logic* 61: 653–86.
- Sher, G. 1997a. "On the Place of Philosophy in Quine's Early Theory", in *The Role of Pragmatics in Contemporary Philosophy*, vol. II, P. Weingartner et al. (eds). Vienna: Wittgenstein Society, pp. 889–95.

- Sher, G. 1997b. "Partially-Ordered (Branching) Generalized Quantifiers: A General Definition". *Journal of Philosophical Logic* 26: 1–43.
- Sher, G. 1998–9. "On the Possibility of a Substantive Theory of Truth". *Synthese* 117: 133–72.
- Sher, G. 1999a. "Is There a Place for Philosophy in Quine's Theory?". *Journal of Philosophy* 96: 491–524.
- Sher, G. 1999b. "Is Logic a Theory of the Obvious?". *European Review of Philosophy* 4: 207–38.
- Sher, G. 2000. "The Logical Roots of Indeterminacy", in *Between Logic and Intuition*, G. Sher and R. Tieszen (eds). Cambridge: Cambridge University Press, pp. 100–23.
- Sher, G. 2001. "The Formal-Structural View of Logical Consequence". *Philosophical Review* 110: 241–61.
- Sher, G. 2002. "Logical Consequence: An Epistemic Outlook". *Monist* 85: 555–79.
- Sher, G. 2003. "A Characterization of Logical Constants Is Possible". *Theoria* 18: 189–97.
- Sher, G. 2004. "In Search of a Substantive Theory of Truth". *Journal of Philosophy* 101:5–36.
- Sher, G. 2005. "Functional Pluralism". *Philosophical Books* 46: 311–30.
- Sher, G. 2007. "Review of *Rationality and Logic*". *Notre Dame Philosophical Reviews* (April): 1–6.
- Sher, G. 2008. "Tarski's Thesis", in *New Essays on Tarski and Philosophy*, D. Patterson (ed.). Oxford: Oxford University Press, pp. 300–39.
- Sher, G. 2010. "Epistemic Friction: Reflections on Knowledge, Truth, and Logic". *Erkenntnis* 72: 151–76.
- Sher, G. 2013. "The Foundational Problem of Logic". *Bulletin of Symbolic Logic* 19: 145–98.
- Sher, G. 2016a. "Truth and Transcendence: Turning the Tables on the Liar Paradox", in *Reflections on the Liar*, B. Armour-Garb (ed.). Oxford: Oxford University Press.
- Sher, G. 2016b. "The Model-Theoretic Argument: From Skepticism to a New Understanding", in *The Brain in a Vat*, S. Goldberg (ed.). Cambridge: Cambridge University Press.
- Sher, G. and Tieszen, R. (eds). 2000. *Between Logic and Intuition*. Cambridge: Cambridge University Press.
- Sider, T. 2011. *Writing the Book of the World*. Oxford: Oxford University Press.
- Siewert, C. 2002/6. "Consciousness and Intentionality", in *Stanford Encyclopedia of Philosophy*. (Fall 2011 ed.), Edward N. Zalta (ed.). <<http://plato.stanford.edu/archives/fall2011/entries/consciousness-intentionality/>>.
- Simchen, O. 2012. *Necessary Intentionality*. Oxford: Oxford University Press.
- Soames, S. 1999. *Understanding Truth*. Oxford: Oxford University Press.
- Sober, E. 1982. "Realism and Independence". *Noûs* 16: 369–85.
- Sorensen, R. 2012. "Vagueness", in *Stanford Encyclopedia of Philosophy*. (Winter 2013 ed.), Edward N. Zalta (ed.). <<http://plato.stanford.edu/archives/win2013/entries/vagueness/>>.
- Sosa, E. 1980a. "The Raft and the Pyramid: Coherence versus Foundations in the Theory of Knowledge", in *Knowledge in Perspective*. Cambridge: Cambridge University Press (1991), pp. 165–91.
- Sosa, E. 1980b. "The Foundations of Foundationalism", in Sosa (1991), pp. 149–64.
- Sosa, E. 1983. "Nature Unmirrored, Epistemology Naturalized", in Sosa (1991), pp. 86–107.
- Sosa, E. 1985. "Knowledge and Intellectual Virtue", in Sosa (1991), pp. 225–44.
- Sosa, E. 1991. *Knowledge in Perspective*. Cambridge: Cambridge University Press.
- Sosa, E. 1994. "Philosophical Scepticism and Epistemic Circularity". *Proceedings of the Aristotelian Society*. Supplementary Vol. 68: 263–90.
- Sosa, E. 1997. "Reflective Knowledge in the Best Circles". *Journal of Philosophy* 94: 410–30.

- Sosa, E. 2009. *Reflective Knowledge*. Oxford: Oxford University Press.
- Spelke, E. et al. 1995. "The Development of Object Perception", in *Visual Cognition*, vol. 2, S. M. Kosslyn and D. N. Asherson (eds). Cambridge, Mass.: MIT Press, pp. 297–330.
- Stanford, K. 2009/13. "Underdetermination of Scientific Theory", in *Stanford Encyclopedia of Philosophy* (Winter 2013 ed.), Edward N. Zalta (ed.). <<http://plato.stanford.edu/archives/win2013/entries/scientific-underdetermination/>>.
- Stanley, J. 2011. *Know How*. Oxford: Oxford University Press.
- Steiner, M. 1998. *The Applicability of Mathematics as a Philosophical Problem*. Cambridge, Mass.: Harvard University Press.
- Stroud, B. 1981. "The Significance of Naturalized Epistemology", in *Naturalizing Epistemology*, H. Kornblith (ed.). Cambridge, Mass.: MIT Press (1985), pp. 71–89.
- Stroud, B. 1984. *The Significance of Philosophical Scepticism*. Oxford: Oxford University Press.
- Suppes, P. 1978. "The Plurality of Science". *Philosophy of Science* 2: 3–16.
- Swoyer, C. 2003. "Relativism", in *Stanford Encyclopedia of Philosophy* (Winter 2010 ed.), Edward N. Zalta (ed.). <<http://plato.stanford.edu/archives/win2010/entries/relativism/>>.
- Tappolet, C. 1997. "Mixed Inferences: A Problem for Pluralism about Truth Predicates". *Analysis* 57: 209–10.
- Tappolet, C. 2000. "Truth Pluralism and Many-Valued Logics: A Reply to Beall". *Philosophical Quarterly* 50: 382–5.
- Tarski, A. 1931. "On Definable Sets of Real Numbers", in *Logic, Semantics, Metamathematics*. Indianapolis: Hackett (1983), pp. 110–42.
- Tarski, A. 1933. "The Concept of Truth in Formalized Languages", in Tarski (1983), pp. 152–278.
- Tarski, A. 1936a. "On the Concept of Logical Consequence", in Tarski (1983), pp. 409–20.
- Tarski, A. 1936b. "The Establishment of Scientific Semantics", in Tarski (1983), pp. 401–8.
- Tarski, A. 1944a. "A Philosophical Letter of Alfred Tarski". *Journal of Philosophy* 84 (1987): 28–32.
- Tarski, A. 1944b. "The Semantic Conception of Truth and The Foundations of Semantics". *Philosophy and Phenomenological Research* 4: 341–76.
- Tarski, A. 1966. "What Are Logical Notions?" *History and Philosophy of Logic* 7 (1986): 143–54.
- Tarski, A. 1983. *Logic, Semantics, Metamathematics*. Indianapolis: Hackett.
- Thagard, P. 1992. *Conceptual Revolutions*. Princeton: Princeton University Press.
- Thalos, M. 2013. "The Government of Logic and Mathematics: A Dialectic". Manuscript, University of Utah.
- Tolley, C. 2006. "Kant on the Nature of Logical Laws". *Philosophical Topics* 34: 371–407.
- Tolley, C. 2013. "Kant on the Generality of Logic". *Proceedings of the 11th International Kant Congress*. Berlin: de Gruyter.
- Väänänen, J. 1985. "Set-Theoretic Definability of Logics", in *Model-Theoretic Logics*. J. Barwise and S. Feferman (eds). New York: Springer-Verlag, pp. 599–643.
- Van Benthem, J. 1983. "Determiners and Logic". *Linguistics and Philosophy* 6: 447–78.
- Van Benthem, J. 2002. "Logical Constants: The Variable Fortunes of an Elusive Notion", in *Reflections on the Foundations of Mathematics*, W. Sieg et al. (eds). Urbana: Association for Symbolic Logic, pp. 420–40.
- Van Cleve, J. 1999. *Problems from Kant*. Oxford: Oxford University Press.
- Van Cleve, J. 2005. "Why Coherence is Not Enough: A Defense of Moderate Foundationalism", in *Contemporary Debates in Epistemology*, M. Steup and E. Sosa (eds). Oxford: Blackwell, pp. 168–80.

- Van Dalen, D. and van Atten, M. 2002. "Intuitionism", in *A Companion to Philosophical Logic*, D. Jacquette (ed.). Oxford: Blackwell, pp. 513–30.
- Vanzo, A. 2010. "Kant on the Nominal Definition of Truth". *Kant-Studien* 101: 147–66.
- Varzi, A. 2003/9. "Mereology", in *Stanford Encyclopedia of Philosophy* (Winter 2015 ed.), Edward N. Zalta (ed.). <<http://plato.stanford.edu/archives/win2015/entries/mereology/>>.
- Wagner, E. H. 1960. "The Unreasonable Effectiveness of Mathematics in the Natural Sciences". *Communications on Pure and Applied Mathematics* 13: 1–14.
- Wang, H. 1987. *Reflections on Kurt Gödel*. Cambridge, Mass.: MIT Press.
- Watson, J. 1968. *The Double Helix*. New York: Mentor.
- Westerståhl, D. 1985. "Logical Constants in Quantifier Languages". *Linguistics and Philosophy* 8: 387–413.
- Whitehead, A. N. and Russell, B. 1910–3/1925–7. *Principia Mathematica*, vols. I–III. Cambridge: Cambridge University Press.
- Wigner, E. 1960. "The Unreasonable Effectiveness of Mathematics in the Natural Sciences". *Communications in Pure and Applied Mathematics* 13: 1–14.
- Williams, B. 2002. *Truth and Truthfulness*. Princeton: Princeton University Press.
- Williams, M. 2001. *Problems of Knowledge*. Oxford: Oxford University Press.
- Williamson, T. 2013. "How Deep is the Distinction between A Priori and A Posteriori Knowledge?", in *The A Priori in Philosophy*, A. Casullo and J. Throw (eds). Oxford: Oxford University Press, pp. 291–309.
- Wittgenstein, L. 1921. *Tractatus Logico-Philosophicus*. London: Routledge and Kegan Paul, 1961.
- Wittgenstein, L. 1953/58/63. *Philosophical Investigations*. Oxford: Basil Blackwell (1972).
- Woodward, J. 2000. "Explanation and Invariance in the Special Sciences". *British Journal for the Philosophy of Science* 51: 197–254.
- Wright, C. 1992. *Truth and Objectivity*. Cambridge, Mass.: Harvard University Press.
- Wright, C. 1999. "Truth: A Traditional Debate Reviewed", in *Truth*, S. Blackburn and K. Simmons (eds). Oxford: Oxford University Press, pp. 203–38.
- Wright, C. 2001. "On Basic Logical Knowledge: Reflections on Paul Boghossian's 'How Are Objective Epistemic Reasons Possible?'". *Philosophical Studies* 106: 41–85.
- Wright, C. 2013. "A Plurality of Pluralisms", in *Truth and Pluralism*, N. J. L. I. Pedersen and C. D. Wright (eds). Oxford: Oxford University Press, pp. 123–53.
- Wright, C. D. 2005. "On the Functionalization of Pluralist Approaches to Truth". *Synthese* 145: 1–28.
- Wright, C. D. 2010. "Truth, Ramsification, and the Pluralist's Revenge". *Australasian Journal of Philosophy* 88: 265–83.
- Wright, C. D. 2012. "Is Pluralism about Truth Inherently Unstable?". *Philosophical Studies* 159: 89–105.
- Wright, C. D. 2015. "Explanation, Truth, Minimalism". Manuscript, CSU Long Beach.
- Yablo, S. 2001. "Go Figure: A Path Through Fictionalism". *Midwest Studies in Philosophy* 25: 72–102.
- Yablo, S. 2014. *Aboutness*. Princeton: Princeton University Press.
- Yi, Byeong-Uk. 1998. "Numbers and Relations". *Erkenntnis* 49: 93–113.
- Yi, Byeong-Uk. 1999. "Is Two a Property?". *Journal of Philosophy* 96: 163–90.

Index

- aboutness 94, 163, 164n4
- absoluteness 307–10
- abstract ix, 20, 55, 67, 79–80, 86, 108, 135–9, 143n27, 149–50, 152, 172, 191–2, 197, 205, 215, 222, 230, 249, 291, 330
 - features ix–x, 73, 77–84, 90–1, 124, 126, 194, 207, 214–15, 256, 272, 292n63
 - knowledge ix–x, 16, 36, 39–40, 51, 56, 73–4, 79–83, 93, 100, 104–5, 109–12, 124, 143, 150–1, 177, 192, 196–7, 204, 208n76, 215, 217, 245, 249–50, 261, 294, 331
 - objects 77–9, 81–3, 104n3, 108, 192, 205, 211–12, 214, 290
- abstraction x, 13, 55, 84, 144, 147, 150, 168–9, 186, 190, 205n70, 209, 218, 240, 278, 291, 325
- abstractness xii, 81, 84n15, 197, 205, 207, 251, 288, 291, 294
 - of logic, *see* logic, abstractness of
- algebra 203, 207–8, 321–2n108, 335n119, 336
- Allison, H. E. 53n21
- Almog, J. 258
- Alston, W. 23, 30, 184n27
- analytic 5, 6n5, 40–55, 58n27, 91, 100–3, 111–17, 126–7, 131, 137, 160, 211, 226, 240, 256, 270–1, 327
- analytic-synthetic dichotomy 40–55, 62, 68, 73, 100–3, 113–16
- analytic-synthetic thesis 41, 42n3, 44, 103, 126
- analyticity 6n5, 49–53, 59n27, 101, 113n16, 115, 127, 211n81, 270–1, 288n57, 327
- Anderson, P. W. 143n29, 144
- aposteriori 42n3, 50, 87–90, 116, 137n20
- application (applicability) problem 80, 108, 155, 203–5, 209, 212, 249, 271, 282, 289, 291–2, 323, 326, 337
- apriori 6n5, 14, 42n3, 47n14, 49–50, 59n28, 67, 71, 73–4, 76, 84, 87–91, 113, 115–16, 123–7, 137, 160, 211, 239, 286n55, 288, 293, 294n, 302n74, 341
 - quasi- 59n28, 89, 113, 116, 211, 288, 293–4, 302
- apriorism 73–4, 84, 90–1, 113, 127
- apriority 87–9, 126, 211n81, 286n55, 288, 293, 302n74, 342
 - quasi- 89, 113, 288, 293–4, 302
- Aquinas, T. 186
- Aristotle 5, 15, 51, 84n15, 106, 186, 208–10, 242
- arithmetic xii, 60, 69, 72, 86, 92–3, 96, 151, 167, 188, 196–204, 207–9, 254, 270, 309n85, 321n108, 322, 324–5
- Austin, J. L. 142, 186
- Azzouni, J. 74
- Baghramian, M. 97n41
- Balaguer, M. 79, 81n9
- Baker, G. P. 150n44
- Bar-Elli, G. 211n81
- Barnard, R. 213n85
- Barrett, R. B. 101
- Barwise, J. 298–9, 308
- basic cognitive situation xi, 74–5, 162–3, 168, 174, 189–90
- basic epistemic situation vii, xiii, 75–6, 91, 120–1, 124, 162, 172
- basic-knowledge predicament 19–20, 244, 250
- basic realism 56, 59, 73, 91–2, 97–8, 127, 162, 207
- basicness xiii, 35, 251, 288, 291–2
 - of logic, *see* logic, basicness of
- Bealer, G. 90, 113
- Beall, J. C. 334–6
- Beck, A. 278n43
- Belnap, N. D. 31–2, 232, 248, 302
- Ben-Menachem, Y. 150n39
- Benacerraf, P. 20, 203, 206–7
- bifurcation
 - analytic-synthetic, *see* analytic-synthetic dichotomy; analytic-synthetic thesis
 - center-periphery, *see* center-periphery bifurcation
 - problem of 80–1, 249
 - of standards of knowledge 40–2, 46n12, 50, 73, 80–1, 116, 328
- bivalence 70, 94–5, 221, 233n116, 239, 261, 275–6, 316–17, 328, 337
- Blackburn, S. 135, 229
- Boghossian, P. A. 46n13, 55n24, 113, 244
- BonJour, L. 22–3, 74, 90, 113
- Bonnay, D. 298n70, 315n96, 331
- Boolos, G. 318
- Boyd, R. 95n40
- Brandom, R. B. 86n21, 150n40, 179, 227n109, 258
- Bueno, O. 302n74
- Burge, T. 241n5
- Burgess, J. 77–8, 308

- cardinality x, xii, 57, 83–4, 92–3, 117, 193–205, 209–10, 213, 272, 287n56, 289, 291, 293, 304, 310, 326
- Carnap, R. 14, 17, 43–6, 50, 52, 111, 168, 246, 327, 334–6
- Cartwright, N. 143–4, 203n63
- Casullo, A. 90n28
- Cavaillès, J. 5
- center ix–x, 15, 40, 42, 49, 55–9, 62, 65, 68–73, 95, 100, 111, 125–7, 202, 207, 236, 259, 270, 272, 340
- center-periphery bifurcation 40–9, 55, 59, 62, 68, 91, 340–2
- center-periphery thesis 41, 55; *see also* CP
- certainty xiii, 19, 121, 126, 242, 250, 288, 293–4, 320
- change (movement)
 - contextual 62, 66, 68, 70–1, 126
 - developmental 68–9
 - temporal 62, 68–71
- Chateaubriand, O. 93n33, 258
- Chomsky, N. 113, 145
- circularity 23, 27, 30–4, 36, 181–2, 243–4, 248, 251, 302n74, 323, 342
 - constructive 27, 30–2, 34, 36, 323
 - destructive 23, 27, 30, 34, 251
 - trivializing 30, 32
- classical logic 52n19, 258, 261, 337
- cognition ix, xi–xii, 8, 10–11, 25, 55, 65, 76, 84–5, 87, 88, 97, 134, 135, 137, 138, 163n3, 164, 166–8, 182–5, 189, 192, 251
- cognitive vii, x–xiii, 3, 10, 13–15, 19, 28–9, 31, 34, 52, 54, 56, 62, 66, 74–6, 80–2, 84–7, 91, 95, 97–8, 120–1, 142, 144, 148, 162–3, 167–8, 171–4, 180n23, 185–6, 189–190, 200–1, 207, 217, 233, 235, 243, 254, 256, 258–9, 295, 323, 341; *see also* basic cognitive situation
 - access problem 80, 95, 203, 206–7
 - route 29, 55, 71, 75, 82, 101, 124, 174, 252, 341
- Cohen, S. 30
- coherentism viii, 21–3, 26–8, 55, 125, 155, 189, 342; *see also* theory of truth, coherence
- command 332
- common denominator 139–152
 - myth of, *see* myth of the common denominator
- common sense 20, 24, 34, 71, 77–8, 82–3, 131–2, 147, 159, 188, 192n39, 194, 211, 245–6, 250, 260–1, 272, 306
- composite correspondence, *see* correspondence, composite
- composite reference 200–2, 217
- compositionality 67–8
- Comte, A. 144
- concept(s) 4–8, 10, 16, 19, 29, 32, 44, 47n14, 49, 51–4, 57, 59, 70–1, 76, 91, 104, 109, 115–16, 125–6, 132, 137, 141–4, 149–50, 153–4, 156–7, 160, 163, 164n5, 166, 168–9, 173, 179, 219–21, 227, 242, 249, 255, 257, 284, 296, 306–8, 334
 - fixity, non-fixity of 53–4, 57, 115, 228n111
- confirmation 32, 51, 57
- confirmation holism 26n10, 114
- consequence 51, 77, 94, 107, 120, 160, 215, 222, 224, 227, 254, 258, 260, 262–9, 273–4, 278, 283, 287, 291, 295, 296–8, 300, 303–305, 313–14, 319, 326, 338, 341
 - biological 227
 - logical xiii, 35, 51n18, 82, 147, 167, 219, 225–7, 234n119, 248, 253, 257–8, 262–74, 277–8, 280–3, 287, 291, 296–9, 302–4, 313, 315, 318–19, 321–2, 328, 335–6, 338
 - material 263, 266–8, 328
 - modal 227
 - nommic 263, 266–8
 - physical 227
- conservative 212, 248
- content ix, 7–8, 12, 41, 47, 53–4, 96, 110, 112, 135, 138, 141, 149, 158–61, 163, 165, 171, 178, 183, 219, 226, 240, 245, 259, 276, 278n42, 295–6, 298
- context x, 21n3, 32, 59, 62, 66, 68, 70–1, 86, 91, 121, 124, 126, 132, 136–7, 142–3, 153, 165, 167, 174, 180, 191n35, 193–5, 204n69, 210, 214–15, 243, 264n, 278n44, 284n51, 286, 291, 295, 304, 310–11, 314, 316, 319, 324n, 326–7, 330n114, 332–3, 337–8
 - of assessment 180
 - of command 332
- continuity through change 71
- Continuum Hypothesis 305, 308, 313
- convention ix, x, xii, 11, 18, 20, 40, 42–6, 51, 53–5, 57, 62, 65–6, 69, 72, 96, 102, 104, 105, 107, 111, 168, 175, 188, 193, 196, 203–4, 215, 241, 246–50, 255, 258, 271, 283, 296
- conventionalism, conventionality 20, 51, 64, 109n11, 168, 189n32, 196, 245–50, 259, 303, 307, 320, 335n, 342
- Copernican revolution 15, 47, 64, 86, 167, 242
- correctness x, xii, 17, 24n6, 52, 58, 96, 107, 151, 162, 172, 188, 245, 247, 258, 295, 334–5
- correspondence x, xii, 60, 63, 66n6, 81, 94, 96, 97–8, 147–8, 152–7, 159, 181–3, 185–93, 198, 200–1, 203–4, 207, 209, 212–18, 220–1, 233–6, 264, 271, 300–1, 336, 340, 342

- composite xii, 66n7, 199–203, 211–13, 215n88, 217, 235
- copy 186, 189, 191–2
- direct, indirect 198, 200, 202, 213–15
- isomorphism 186, 189
- logical 216, 236, 301
- manifold xii, 186–7, 189, 192, 208, 216–19, 235, 340, 342
- mathematical 60, 187, 189, 192–4, 199–201, 203, 205–6, 208, 212, 216–18
- mirror 66, 148, 152, 186, 189, 191–2
- picture 66, 152, 186, 191
- route of 189, 191, 194, 199, 206–7, 216, 235, 340
- traditional, non-traditional vii, x, xii, 36, 186–7, 189, 191–2, 201, 206–7, 235
- CP 41–3, 46–50, 55–56, 105, 112; *see also* center-periphery thesis
- criterion
 - of formality 284, 286, 297–8, 305
 - invariance 297–9, 302–3, 305–14, 322, 331
 - Kantian 135–8
 - for logical constants 234–5, 253, 287, 296–7, 299, 303–4, 312–13
 - for logical operators 234n119, 279, 286–7, 298, 303–5, 308, 312–13
 - of logicity 286, 298n70, 302, 305–14, 319n106, 322, 331, 333
 - myth of the 135–142
 - of ontological commitment, *see* ontological commitment
 - of truth 135–8, 186
- Daniels, N. 32n17, 34n19
- Darwin, C. 51, 106, 228
- David, M. 133, 187
- Davidson, D. 103, 112, 133, 153, 182–4
- definition
 - of denotation 226
 - of formula 223
 - inductive 223
 - recursive 32, 217, 223–4, 226, 230n113
 - of satisfaction 224, 226
 - of truth 32, 147, 217, 220–33, 279n45, 281
- deflationism xi, 12, 105n4, 113, 132–3, 142, 146, 148, 342
- deflationist x–xi, 21, 39, 105, 113, 127, 131–3, 136–8, 145, 148, 154, 158, 174, 203, 260, 336
- denotation 225–6, 228, 257, 280–2, 298, 300
- Descartes, R. 17, 70
- Devitt, M. 133, 153
- Dewey, J. 86n21, 182
- diallelon 181–5
- Dilman, I. 101
- discovery viii, x, xiii, 12, 15, 30–1, 45, 60, 69, 78, 87, 88n23, 100–1, 113, 127, 131–2, 147, 151–2, 181, 204n69, 211, 249–50, 253, 293, 299, 323, 331, 341
- disquotation 63, 158, 187
- disquotational schema 63, 145n33, 158, 221n97
- disunity xi, 134, 137, 142, 144, 149–50, 153–4
 - challenge 134, 142–4, 146, 148, 150–1, 153–4
 - moderate 142–4, 146, 148, 151
 - radical 134–40, 142, 153
 - of science 143–4
 - of truth, *see* truth, disunity of
- diversity xi, 16, 67–8, 75, 88, 134, 143, 146–51, 160–1, 165–6, 189, 207, 218, 230, 337
 - of truth, *see* truth, diversity of
- DQ 117–22
- Dreben, B. 102–3
- Duhem, P. 117–19
- Duhem-Quine thesis 117–22
- Dummett, M. 25–6, 40–1, 48, 54, 58–9, 68, 73, 92n30, 94–6, 112, 131–3
- Dupré, J. 143–5
- Dutilh Novaes, C. 278n42, 302
- Dworkin, R. 175
- dynamic 31, 33–4, 40, 59n28, 62, 68–9, 71–2, 76, 100, 104, 116, 125–6, 147, 165–7, 180, 207–8, 233, 336, 342
 - model of knowledge ix–x, 35, 37, 57, 59, 61–2, 68, 72, 76, 91, 121, 126, 161–2, 180, 213, 215, 233, 302
- Dyson, F. xi, 147
- Eco, U. 53n22
- Einstein, A. 51, 60, 67, 86, 106, 123–4, 328, 337n123
- Elgin, C. Z. 23
- empirical knowledge, science viii, x–xi, 5, 8, 14, 24, 26, 29, 42, 67, 73–4, 88, 91, 101, 103, 105–6, 108, 117, 121, 124–5, 137, 145, 148, 164, 204, 211–12, 247, 249–51, 259, 293–4, 328, 330–1
- empiricism ix, 10, 18, 36, 48–50, 59, 64n3, 67, 73–4, 77, 80, 84, 100–13, 117n20, 123–7, 157, 195, 203, 207, 213–15, 250, 260, 316, 327, 342
 - Quine's, *see* Quine, Quine's (radical) empiricism
- Enderton, H. B. 223
- epistemic freedom vii–x, xii–xiii, 3, 12–14, 16, 24, 33, 35, 39–40, 48, 58–9, 66, 72–3, 82, 86–7, 98–9, 104, 113n17, 117n20, 121n30, 126, 167, 180, 200, 207, 278n44, 301n73, 336, 340–41; *see also* freedom
 - extended sense 13
 - natural sense 13

- epistemic friction i, iii, viii, ix–x, xiii, 1, 3–6, 8–9, 12–13, 16–17, 24, 35, 39, 48, 72–3, 89, 91, 104, 126, 341
 principles of 9–12
 epistemic methodology 16–17, 21, 39, 99, 340
 epistemology vii, 9n8, 17–21, 25, 43, 46–8, 50, 55, 59, 64, 67, 71–2, 76–7, 80, 90n28, 93, 103, 105, 124, 152, 295; *see also* theory, of knowledge
 naturalized 70, 102–3
 equivalence schema 145–6, 221n97
 error theory 192
 Esfeld, M. 104
 Etchemendy, J. 297n69
 ethics 92, 147–8, 153, 174, 176–7, 217, 295, 341
 evidence 9, 17, 25, 40, 51, 55–6, 66, 100–2, 107, 123, 162, 174, 181, 211, 246, 259, 329
 experience 5–6, 8, 10, 14, 18, 25, 41–6, 50, 52–3, 56, 58n26, 62, 86, 88, 89, 95, 100–2, 108–9, 115, 118–19, 122–3, 127, 137, 211, 246, 327, 329–30
 explanation 9, 12, 19, 25, 35n20, 55, 57, 63, 131–3, 136, 139, 144, 150–1, 157–9, 212, 227, 245, 260, 267, 288, 314
 externalism 18, 24n6

 fact 12, 35n20, 41–2, 50–1, 59, 96, 102, 104, 115–16, 183, 248–9, 265n22, 278–9, 301, 328, 335
 factuality 6, 12, 24n6, 39, 42n3, 45–6, 49, 55–8, 60, 63–5, 67–9, 72, 87–8, 105–7, 109, 116, 143, 168, 187, 211, 248–9, 259, 261–2, 267, 271, 290, 327, 329, 331
 family resemblance 53, 150
 feature
 abstract, *see* abstract, features
 formal 22, 83, 91, 109, 194–6, 198, 202, 204–12, 271–4, 277, 283, 291–2, 300, 324, 326, 328
 mathematical 193–6, 207–9
 physical 82, 84, 90–1, 195, 272
 Feferman, S. 299, 302–8, 310–12, 314–15, 331
 Fermat's last theorem 86, 329
 Feynman, R. 194
 fictionalism 192, 208, 210n80, 212–13
 Field, H. 132–3, 146, 152n45, 158, 172, 212, 220n95, 229–30, 295, 318–19, 327
 figuring out x, 9, 81, 84–7, 134, 173, 192–3, 199–201, 207n74, 211, 234, 249n13, 254, 342
 Fine, K. 35n20, 73, 188, 208n77, 216n88, 249n12
 Finkelstein, D. 52n19, 69
 first-order
 logic 31, 106, 164, 202, 225, 254, 277, 287, 298–9, 305, 309, 312–13, 319, 322, 333
 thesis 294, 298–9
 fixed x, 53–4, 56–8, 68, 70–1, 115, 161, 227–8, 231, 267, 280–1, 323–4, 333
 fixing and unfixing 70, 180
 flexible realism, *see* realism, flexible
 Floyd, J. 150n41
 Fodor, J. 25–6, 143n29, 145
 formal
 adequacy 220–2
 facet of reality xii, xiii, 319, 321
 feature, *see* feature, formal
 impossibility 274
 law 57, 196–8, 200, 205–6, 269, 271–4, 279–83, 287–93, 296, 301–2, 307, 316–21, 326, 328, 330–3
 operator 274n36, 276–81, 287, 297, 313, 324n111, 325
 property x, 83, 195, 203, 272, 276
 structure 21, 23, 52, 57–8, 67, 94, 272, 274–6, 279, 282, 284, 287, 304, 307, 310, 313–14, 317n100, 322–3, 325, 328–31, 337
 truth, *see* truth, formal
 formality xii, 5, 83n14, 205–6, 239, 269n29, 271, 274–9, 282–8, 291–4, 297–300, 305–6, 308, 312–17, 321, 324–5, 331, 334
 of logic, *see* logic, formality of
 thesis 271, 274, 279, 283, 287–8, 294, 297–300
 general theoretical account 274, 276–83
 joint perspective 287–8
 specialized mathematical account 283–7
 formula 223–6, 281, 308
 foundation
 of knowledge 8, 17–18, 20, 23–4, 26, 30, 33, 35, 240, 243
 of logic vii, xiii, 16, 32n16, 36, 86, 107, 235, 239–338, 340
 without foundationalism viii, 21, 22–4, 26, 28, 243
 foundational holism viii–ix, 20, 24–30, 34–5, 59, 104, 161, 202, 215, 244, 246, 250–3, 267n26, 273, 302, 309, 320–1, 337, 342
 foundational methodology viii, xii, 17, 20–1, 23–4, 27n13, 29, 125–6, 243, 246, 253, 302, 309, 321
 foundational problem of logic xiii, 235, 239–53
 foundational project viii–ix, 17, 21–30, 71, 73, 86, 240, 243, 250, 305, 340
 foundationalism viii–ix, 17–24, 27–9, 35, 71, 86, 101, 124–6, 181, 243–4, 246–50, 302, 307, 320–1, 323, 325, 327, 342
 illusion of 17
 strict ordering requirement 18–24
 Fraenkel, A. A. 275, 310
 freedom 15–16, 23, 25, 45, 174, 188, 251, 332–3, 336n122, 340–2
 active 14–15, 40, 72

- epistemic, *see* epistemic freedom
- passive 14
- Frege, G. 15, 41n3, 69, 79, 96, 164, 180n23, 197n50, 208, 210–11, 215n88, 219n94, 241, 259, 278–9, 295–6, 320
- Freud, S. 173
- friction vii, xii–xiii, 4–13, 16, 19, 24, 29–30, 35, 59–60, 111, 127, 131, 174, 181, 236, 242, 247, 250–1, 270, 303, 329, 340–1
 - epistemic, *see* epistemic friction
- frictionless ix, xi, 6–7, 30, 80, 159, 316
 - theorizing 4–6, 80, 338
- Friedman, M. 59, 73, 90–1, 113–14, 117, 122–7, 246
- front of knowledge, *see* knowledge, front of
- fruitful balance xi, 147, 160–1
- functional approach 253, 255n1, 322, 326
- functionalism 154
- fundamental principle of truth xi, 162, 171–9, 183–5, 188–9, 217–19, 233
 - immanence xi, 64, 72, 161–79, 183–5, 188–9, 217–18, 233
 - normativity xi, 65n5, 147, 161, 169–72, 218, 233, 236, 239, 288, 292, 294–6
 - transcendence xi, 72–3, 135, 137, 157, 161, 164–71, 178–89, 202, 204, 213, 232–3, 258–9
- fuzzy logic 261, 337
- generality xi–xiii, 5, 32, 65, 71, 136, 142–5, 150, 153, 157–60, 194, 197, 198, 206–7, 229–30, 236, 255–6, 271, 282, 287–8, 290–2, 294, 298–300, 302, 310, 313–17, 323, 326, 330–1, 334–5, 337
 - of logic, *see* logic, generality of
- genetic fallacy 52, 54, 69, 270n32
- geometry 51, 60, 69, 119n28, 123, 208–9, 258, 295, 322, 336
- Gibbard, A. 92n31, 176
- Gibson, R. F. 46n13, 100–1, 103–4
- Glymour, C. 25–6, 31–2, 112
- Gochet, P. 46n13, 101
- God's eye view xi, 98, 166–7, 181, 183, 342
- Gödel, K. 31, 69–70, 85, 109, 151, 167–8, 241, 249, 284, 299, 307–8
- Goldfarb, W. 259
- Goldman, A. 131–2
- Gómez-Torrente, M. 302–3
- Goodman, N. 31–32, 182, 191
- Grayling, A. C. 95n40
- Grice, H. P. 46n13, 55n24, 113–16
- grounding
 - of knowledge viii, 8–11, 16, 17–36, 39–40, 90, 92, 98, 104, 125, 143, 340, 342
 - of logic ix, xii–xiii, 35–6, 39, 60, 67, 74, 90, 95, 105, 107, 219, 235, 240–52, 255–71, 305, 308, 310, 315–16, 319–21, 326, 333, 337, 340–1
- in the mind ix, 8–9, 11, 17–36, 40, 60, 104, 107, 186n30, 219, 255, 337
 - a note on 35
 - in reality/world viii–ix, xii–xiii, 7–12, 16–36, 39–40, 60, 67, 74, 80, 90–2, 95, 98, 104–5, 107, 111–12, 125–7, 143, 162, 186n30, 193, 219, 255–71, 310, 315–16, 319, 326, 334, 337, 340–1
- Grünbaum, A. 114, 117–22
- Gupta, A. 31–2, 113, 133, 232
- Haack, S. 23
- Hacker, P. M. S. 150n44
- Hacking, I. 143n29, 145
- Hale, B. 320
- Hanna, R. 14n11, 112–13, 136–7, 239, 240n3, 256, 322n109
- Hanson, W. H. 294n68, 302–3
- Harding, S. G. 118–19
- Harman, G. 65n5, 295
- Hawking, S. W. 145
- Hawthorne, J. 73, 90
- Hellman, G. 24, 208n77, 215n88, 305n77, 325
- Hempel, C. 157–8, 182
- Henkin, L. 31, 68n9
- Herburt, G. K. 119n27
- Hesse, M. 71n12
- HG-transcendence 166–7
- HH-transcendence 166–7
- hierarchy 18, 20, 33, 57, 71, 149, 157, 166n9, 184, 188, 202, 229, 231–4, 244–5, 319
 - Tarskian, *see* Tarski, Tarskian hierarchy
- Higginbotham, J. 299
- Hintikka, J. 68n9
- Hodes, H. 202n62, 210n80, 215n88
- Hofstadter, A. 25n8, 26n10, 112
- Hofweber, T. 200n57, 202n62
- holism viii–ix, xiii, 25–30, 33, 40–1, 48, 50, 55, 59n28, 67–8, 71–3, 76, 86, 100, 111, 114, 117, 122–7, 143, 149, 152, 157, 167, 179–80, 182, 194, 203, 207–8, 211, 234, 253, 255, 291–2, 294, 302n74, 323, 325, 331, 340
 - confirmation 26n10, 114
 - foundational, *see* foundational holism
 - methodology, *see* methodology, foundational-holistic
 - one-unit 25–6, 100, 112
 - Quine's, *see* Quine, Quine's holism
 - relational 26, 101, 111, 115, 117–18, 120–1
 - structural 26, 55, 121, 124
- Hookway, C. 101, 103
- Horgan, T. 208, 213–15
- Horwich, P. 132–3, 145, 174n17
- Hughes, R. I. G. 52n19

Hume, D. 15, 20, 86, 172–3, 182

Hylton, P. 46n13, 101, 103–4

idealism 50, 55, 73, 97, 157–8, 341

identity problem 203–4

imagination vii, 5, 75, 173, 175, 189, 200–1

immanence, *see* fundamental principle of truth, immanence

of knowledge 64, 72

immanent

mode of thought 163, 168, 171

transcendence 72, 168n11

incommensurability 71, 180, 217, 302

independence 91, 95–8, 190

indeterminacy 40, 113, 121, 341

indispensability 108, 122, 211

Infeld, L. 337n123

inference

logical 15, 216–17, 240, 243, 254, 258, 262, 279, 283, 330–2, 334

method of xiii, 218–19, 255, 262, 294, 315, 323, 326, 333–4

mixed 156, 216, 249

theory of 290, 314–15, 323, 338

valid 31n15, 65, 70, 117, 156, 240, 271, 280, 287, 333

infinite regress 23, 32–3, 181–2, 232, 234

instrumentalism 104, 108–9

intellect vii, ix–x, xiii, 73–4, 76, 78, 84–91, 99–101, 103, 113, 125, 127, 131, 149, 173, 185–6, 189–90, 198, 200–1, 207, 210–11, 242–3, 293–4, 331, 336, 341; *see also* reason

intentionality 163

interface

with the mind ix, 40, 62, 98

with reality (the world) ix, 40, 62, 98, 111, 236

internal realism, *see* realism, internal

internalism 18, 24n6

intuition 5, 7–8, 10, 14, 18, 20, 84–5, 239, 245, 249–50, 252, 283, 306

mathematical 84–5

pure 10, 14, 20, 84, 245, 249

rational 18, 249

intuitionism 65, 107, 320, 334, 337

intuitionistic logic 334, 337

invariance xiii, 71, 89, 196n49, 205, 207, 209, 276, 278, 285, 288–90, 292–4, 311, 315–16, 322n108, 330

criterion, *see* criterion, invariance

degree of xiii, 89, 198, 205–6, 288–91, 316–17, 329, 334

under homomorphisms 331

under isomorphisms 285–6, 288, 297–8, 302, 305–7, 310, 312–13, 315, 331

under permutations 298n70

invariance under isomorphisms
criterion 298–314, 322

necessary criterion 307

necessary and sufficient criterion 312–14

objections to 302–6, 308, 310, 314

overgeneration of 303–4

sufficient criterion 305, 312–13

undergeneration of 303

isomorphism 186, 189, 207, 284–6, 288, 292, 297–8, 302, 305–8, 310, 312–15, 322, 324–5

James, W. 149–150, 164, 182

Jech, T. J. 308n84

Jenkins, C. S. 91, 95n40

justification viii, x, xiii, 7, 9, 12, 17, 19, 24–5, 30, 32, 34, 39–41, 53n20, 55, 57, 65, 71, 76, 78, 84n17, 87, 94, 96, 100, 104–9, 113, 116–17, 127, 131, 148, 151, 154–5, 162, 169, 171n, 174, 181, 212, 244–6, 251, 253, 259, 331, 338

Kant, I. 4–15, 17, 45, 47, 52–3, 64–6, 70, 73, 80, 83–9, 93, 135–8, 142, 148–9, 162, 164, 166–7, 172, 182–3, 241–3, 250, 256, 259, 277, 341

Keenan, E. L. 300

Keisler, H. J. 299, 309n86

Khlentzos, D. 95n40

Kim, J. 113

Kirkham, R. L. 231

Kitcher, P. 157, 184, 191n36

knowledge iii, vii–viii, ix–x, xii, 3–7, 9–26, 28–30, 32–3, 35–6, 39–44, 47–8, 50–7, 59–60, 62, 64–8, 70, 72–8, 80–2, 84–98, 100–5, 107, 109–12, 116–18, 120–7, 131, 135–8, 141, 143–5, 147–8, 151–2, 157, 160–3, 165–9, 173–6, 179n21, 181, 183–4, 187–9, 192, 195, 203, 207, 210–13, 215–16, 218, 229, 233–6, 239, 242, 244–52, 254–5, 259–60, 262, 268, 271–2, 278–80, 284, 289–93, 295, 300, 303, 307, 315, 319n104, 323, 326–37, 331, 333–6, 338, 340–1

abstract, *see* abstract, knowledge

basic 18–20, 244, 250

predicament, *see* basic-knowledge predicament

dynamic model of, *see* dynamic, model of knowledge

empirical, *see* empirical knowledge, science

foundation of, *see* foundation, of knowledge

front of 52, 62, 68–70

grounding of, *see* grounding, of knowledge

intellectual x, 249n13

line of 87–9

logical vii, ix, 73, 246, 248, 251

mathematical 6, 73, 81, 89, 91, 108, 193, 207, 242, 322

model of, *see* model of knowledge

- philosophical 151, 303
- physical 212
- quasi-apriori, *see* apriori, quasi-
- Quinean model of ix, 39–61, 72, 86, 99, 100–27
- rear of 58, 62, 68–70
- scientific 8, 66, 86, 101–2, 109–10, 125–6, 167, 211, 242, 246
- system of, *see* system of knowledge
- theoretical viii, 3n1, 20, 24, 100, 177, 254
- theory of, *see* theory, of knowledge
- Korsgaard, C. 169, 172
- Kreisel, G. 318–19
- Kripke, S. 89, 195n46, 231–3
- Kuhn, T. 66, 69, 71, 120, 182, 204
- Lammenranta, M. 30
- language
 - as an arbiter of ontology 74, 79–81
 - formalized 219, 221–2, 225, 229, 231, 257
 - meta 32, 166–7, 178–9, 183, 221–3, 231–2
 - natural 74, 81, 175, 210, 227, 231–3, 239, 252, 302–4, 319
 - object- 32, 164, 187, 222, 231–2, 239n2
- large ontology problem 203, 206
- Laughlin, R. B. 143n29, 144–5
- law
 - empirical 91
 - formal 57, 196–200, 205–6, 269, 271–4, 279–80, 282–3, 287–93, 296, 301, 307, 316–17, 319, 321, 326, 328, 330–3, 337, 342
 - logical 5, 51, 57, 63, 65, 89, 106–7, 247–8, 251, 259–61, 272, 289–92, 296, 316–17, 328, 330, 333, 342
 - mathematical 194n43, 200, 204–9, 325
 - of nature 51, 110, 167, 263, 268, 328
 - physical 58, 89, 123, 194n43, 197, 205, 240, 295, 301, 317, 328
- Lear, J. 208–9
- Leibniz's law 261
- Lepore, E. 25–6
- level 32, 43, 92–3, 145–6, 164, 167–8, 176, 178, 184–5, 188, 192–3, 195, 199–204, 210, 213–14, 224, 227, 231, 234, 264–9, 275–7, 282, 285–6, 290, 298, 304, 306n81, 309, 321, 324
- Lévy, A. 310, 319n105
- Liar paradox 31, 151, 221, 231–3
- Lindström, P. 167, 286, 299, 311
- line of knowledge, *see* knowledge, line of
- linguistic turn 103, 157
- linguistics 14, 68, 133n8, 254, 284, 294, 299
- Linnebo, Ø. 78–9
- logic
 - abstractness of xii, 81, 205, 251, 288, 291, 294
 - basicness of xiii, 35, 251, 288, 291–2
 - classical, *see* classical logic
 - constrained by the world 261–2, 267
 - enabled by the world 262, 267
 - error in 5, 240, 242, 244, 248, 251, 261–2, 293, 327–8, 330
 - formality of 239, 278n43, 291, 293, 297, 321
 - foundational problem of xiii, 235, 239, 241, 243, 252–3
 - function (role, task) of 219, 253, 258, 280, 253, 255, 326
 - fuzzy, *see* fuzzy logic
 - generality of 255–6, 271, 287–92, 294, 298–300, 302, 313–17, 323, 326, 330–1, 334–5, 337
 - grounding of ix, xii–xiii, 35–6, 39, 60, 67, 74, 90, 95, 105, 107, 219, 235, 240–52, 255–71, 305, 308, 310, 315–16, 319, 320–1, 326, 333, 337, 340, 342
 - intuitionistic, *see* intuitionistic logic
 - laws of, *see* law, logical
 - mathematical, *see* mathematical logic
 - and mathematics vii, ix–x, xiii, 21n3, 31–2, 36, 59–60, 67, 69, 77, 81n10, 87n22, 89–90, 95, 109, 117n20, 123–4, 177, 210–11, 215, 236, 241, 247, 249–50, 273, 305–6, 320–3, 325–6
 - meta, *see* metalogic
 - modal, *see* modal logic
 - modal force of xiii, 12, 35, 89, 196n47, 236, 255, 260, 262–6, 268–9, 271, 275, 282–3, 287–8, 290–1, 294, 299–300, 302, 313–14, 316–18, 323, 326, 328, 330–1, 334–5, 337; *see also* logic, necessity of
 - necessity of 116n19, 258, 268, 281–2, 288, 291–2, 294, 297–8, 302, 316–19, 327–8, 337; *see also* logic, modal force of
 - normativity of 65n5, 170n12, 236, 239, 294–6
 - obviousness of 245–6, 283, 327
 - paraconsistent, *see* paraconsistent logic
 - predicate 278, 298
 - quantum, *see* quantum logic
 - quasi-apriority of 89, 113, 288, 293–4
 - revision in, *see* revision, in logic
 - scope of 107, 236, 331, 333–4
 - sentential 31, 63, 278–9, 298, 311–12, 314, 333
 - topic neutrality of xiii, 236, 254, 288, 291, 294
- logical consequence xiii, 35, 51n18, 147, 167, 219, 225–7, 234n119, 253, 257–8, 262–74, 277–8, 280–3, 291, 296–9, 302–4, 313, 315, 318, 322, 328, 335–6
- Tarski's definition of, *see* Tarski, Tarski's definition of logical consequence
- logical consistency 12, 328
- logical constants 218–19, 222, 225–30, 234–5, 248, 253, 280–2, 286–7, 290–1, 294–304, 309, 312–14, 319n106, 321, 323, 326, 328, 331, 333, 335, 338

- logical knowledge, *see* knowledge, logical
 logical laws, *see* law, logical
 logical operators 223, 228, 234n119, 278–9,
 281, 287, 291, 293, 297–8, 304–5,
 307–15, 325–6, 330
 logical pluralism 332, 334, 338
 logical positivists 43, 50, 52, 66, 102
 logical structure 12, 146–7, 218–19, 226, 229,
 256, 274, 279, 281–2, 300–1
 logical system *x*, 235, 241, 248, 261–2, 270–1,
 279–80, 284, 286–7, 299–300, 302,
 307n83, 309–11, 313, 322, 326, 329,
 333–4
 logical theory, *see* theory, logical
 logical truth 64, 67, 106–7, 152, 218, 234n119,
 236, 245, 248–9, 253–4, 258, 282–3, 294,
 300–1, 307, 318–19, 321, 328
 logical validity 12, 30, 65, 117, 156, 216, 239n2,
 254, 280, 287, 333
 logicity 286, 302, 305, 308–14, 322, 331, 333
 principle, *see* truth, logicity principle of
 logicism *xiii*, 157, 241, 320
 Longueness, B. 5, 241n7
 look and see 100, 150, 342
 Lynch, M. 133, 153–60, 187
- Maaß, J. G. 53
 Mac Lane, S. 318n103
 MacFarlane, J. 8n6, 180, 278n42, 295–6, 303
 Maddy, P. 239–40, 256, 258, 314n92, 322n109,
 330n114
 Maginot line 52, 104
 manifold correspondence, *see* correspondence,
 manifold
 Marcus, R. B. 112
 material adequacy 220, 222, 224
 mathematical features, *see* feature,
 mathematical
 mathematical instrumentalism 104, 108–9
 mathematical knowledge, *see* knowledge,
 mathematical
 mathematical logic 224–5, 240–1, 261, 281n50,
 284, 287, 299, 317, 333–4, 338
 mathematical proof 89
 mathematical structuralism 24, 321, 324–5
 mathematical truth 60, 75n4, 83n14, 187,
 192–4, 202–5, 208–12, 214, 272, 324, 326
 mathematics
 as algebra problem 203, 207, 322n108, 335–6
 applicability problem 80, 155, 203–5, 209,
 212, 249
 applied 197
 cognitive access problem 80, 203, 206–7
 identity problem, *see* identity problem
 indispensability of 108, 122, 211
 meta, *see* metamathematics
 subject matter of 196–9, 200, 208
 truth in 60, 75n4, 83n14, 187, 192–4, 202–5,
 208–12, 214, 272, 324, 326
 and the world 194–6
 mathematism 320–1
 May, R. 299
 McCarthy, T. 286n55, 302–3
 McDermid, D. 182, 184n27
 McDowell, J. 4, 7–12, 98, 148, 166n8, 341
 McGee, V. 232–3, 298n70, 311, 318–19
 McGinn, C. 109
 meaning 51–4, 68, 71, 102, 112, 115–16, 141,
 148, 267, 304, 327
 metalanguage, *see* language, meta
 metalogic 31, 117, 151, 168, 225, 239–42, 254,
 283, 299
 metamathematics 206n72, 241, 303–4
 metatheory 32, 51, 120, 164, 166, 202, 224,
 249, 322
 metaphysical realism 98, 166, 213
 metaphysics 4, 6, 47n14, 80, 105, 166
 methodology
 foundational-holistic *viii*–*ix*, 20, 24–30,
 34–5, 59, 104, 244, 246, 250–3, 267n26,
 273, 320, 337, 342
 foundationalist *viii*–*ix*, 17–18, 20–1, 23–4,
 27n13, 29, 125–6, 243–4, 246, 250,
 302, 321
 holistic *ix*, 24, 30, 35–6, 39, 210, 213, 233,
 244, 250, 340
 Miller, A. 97n42
 mind-independence 78, 95, 190
 minimalism 12, 39, 127, 132n, 135–6, 336
 mirroring 66, 148, 152, 186, 189, 191–2
 mixed inference 156, 216, 249
 mixed sentence 156, 216, 249
 modal force 192n39, 194–6, 203, 205–7, 273,
 325; *see also* logic, modal force of;
 necessity
 problem 203, 205
 modal logic 332, 338
 modality 112, 317
 mode of thought
 immanent *xi*, 163
 normative *xi*, 169, 171
 transcendent *xi*, 166, 168, 183
 model of knowledge *vii*, *ix*–*x*, *xii*–*xiii*, 3, 16, 30,
 35, 39–40, 47–9, 55–6, 72, 74, 86, 91,
 97–9, 101, 114, 126, 144n30, 161–2,
 177, 180, 213, 215, 233, 235, 272, 302,
 340, 342
 Carnap's 45, 111
 dynamic, *see* dynamic, model of knowledge
 Quine's, *see* Quine, Quine's model of
 knowledge
 static 342
 model theory 136, 269n31, 297, 299, 302n74,
 318, 322, 328

- Moltmann, F. 202n62
Moore, G. H. 53n20
moral skepticism 176–7
morality 65, 177, 217
Morrison, M. 158
Mostowski, A. 278, 286, 299
myth of the common denominator 139,
141–2, 150
myth of the criterion, *see* criterion, myth of the
myth of the given 7, 20

Nagel, T. 178
NAS, *see* no analytic-synthetic bifurcation thesis
natural language 74, 81, 175, 210, 227, 231–3,
239, 252, 302–4, 319
natural science 5, 12, 91, 108, 123–6, 152, 204
naturalism 14, 29, 70, 77, 113, 122–5, 157–8,
164, 214, 240n3, 256, 258, 321
naturalized epistemology, *see* epistemology,
naturalized
NB, *see* non-bifurcation thesis
necessity 77, 89, 195n46, 206, 209, 274n36;
see also logic, necessity of; modal force
Neurath boat 26–9, 32–3, 120n29, 122, 161,
167, 179, 202, 213, 251
Neurath, O. 26n12, 70, 182
new theory of mathematical truth 192–218
Nicolas, D. 93n35
no analytic-synthetic bifurcation thesis
(NAS) 41–59, 62, 66, 68, 91, 100–1,
104–8, 111–14, 116
nomic force 263
nominalism 73–4, 77–8, 82–4, 126, 194–6,
212–13, 259, 302n74, 342
non-bifurcation thesis (NB) 40, 69, 72, 104
non-fixity 53–4, 57, 115, 228
non-veridical discourse 332
normative viii, xi, xiii, 13, 17, 21, 34–6, 46–8,
53–4, 57, 62, 65n5, 71, 76, 111, 115–16,
123–4, 127, 152, 161, 169–71, 189, 219,
233, 236, 239, 242, 245, 288, 292,
294–6, 302
mode of thought xi, 169, 171
normativity xi, 146–7, 161, 169, 171–2, 218,
233, 292, 295–6
of logic 65n5, 170n12, 236, 239, 288, 292–6
Nozick, R. 288

object language 32, 164, 187, 222,
231–2, 239n2
objectual descent 63
observation statement 41, 57, 71, 89, 110
obvious, obviousness xii, 20, 53, 106, 158, 241,
245–6, 250, 283, 327
ontological commitment 40, 82n11, 92–3, 108,
211–12, 306–7
ontological relativity 113

ontological unit 93
ontology 58, 60, 64, 66–7, 72, 74, 79–83, 92,
108, 148, 152, 194, 196–7, 203, 206, 210,
212n82, 214, 265n22, 314, 341
language as an arbiter of, *see* language, as an
arbiter of ontology
parity with language, *see* realism, syntax-
ontology parity
operator
formal 274n36, 276–81, 287, 297, 313,
324n111, 325
logical 223, 228, 234n119, 278–9, 281, 287,
291, 293, 297–8, 304–5, 307–15,
325–6, 330
Oppenheim, P. 158

Pagin, P. 68n10
paraconsistent logic 334
Parsons, C. 24, 85, 250n15, 324
Patterson, D. 220n95
Peacocke, C. 73, 286n55, 302
Pedersen, N. J. L. I. 153–4
perception vii, ix–x, 5, 8, 20, 30n15, 73, 76,
84–5, 87, 91, 104, 125, 151, 189–90, 207,
210, 294
periphery ix–x, xii–xiii, 40–9, 55–60, 62–5,
67–71, 73, 91, 95, 100, 111, 123, 125–7,
143, 177, 202, 207, 236, 259–60, 262,
270, 272, 327, 329, 340–2
Peters, S. 300, 313n91
philosophical theory, *see* theory, philosophical
philosophy viii, x–xi, xiii, 6–8, 11–12, 15, 20–1,
24, 29–31, 39, 47, 49–50, 56, 59, 62,
64–5, 67–72, 77, 85, 93, 98, 103, 109,
111–12, 114–15, 124–5, 127, 131–4, 145,
148–50, 154, 157, 159, 163n2, 165–8,
178, 194, 203, 206–7, 210, 219n94, 226,
235, 240–1, 243, 249n13, 254, 256, 293,
303, 307, 316, 341
physicalism 77, 146, 157, 212, 230n112
physics 54n24, 58, 60, 67, 69, 87n22, 117,
120–1, 123–4, 127, 136, 144, 153, 165,
187, 190, 194, 205, 209, 211, 217, 240,
242, 254, 272, 289, 295–6, 315, 323, 325,
328, 337
Pines, D. 143–5
platitudes 155–61
Plato 4, 79, 96
Platonism 20, 36, 67, 73–4, 77–84, 91, 117n20,
126–7, 192, 195, 203, 205, 207, 215,
249–50, 259, 302n74, 307, 316, 342
pluralism
logical 236, 332, 334–8
truth 151–7, 215–16, 249n14
Popper, K. R. 14
positivists, logical, *see* logical positivists
Posy, C. 85n18

- pragmatism 50, 55, 107–8, 149–50, 189, 243, 249, 303, 341
 Prauss, G. 136n17
 Priest, G. 258
 Prior, A. N. 248
 proof
 logical 22, 259n7
 mathematical 89
 theory 280, 283–4, 323, 331, 333
 property
 formal *x*, 83, 195, 203, 272, 276
 physical *x*, 65, 94, 197, 205, 272, 288–9, 293, 316
 psychology *x*, 14, 29, 34, 59, 67–70, 124, 144, 168, 207, 210, 254, 256
 Putnam, H. 51–4, 73, 97n41, 98, 101, 108–10, 112, 133, 150n40, 166, 182–3, 208, 257

 quantification 68n9, 106, 170, 174, 229
 quantum logic 51–2, 65, 107, 261, 330, 337
 quasi-apriori, *see* apriori, quasi-
 quasi-apriority, *see* apriority, quasi-
 question of truth 169–79, 188, 233, 235, 334–5
 quietism 8, 12, 21, 39, 127, 132n2, 341
 Quine, W. V. vii, ix–*x*, 14, 20–1, 25–7, 40–55, 57, 59–64, 66–7, 69–73, 82n11, 92–3, 99–118, 121–7, 163–4, 182, 208, 211, 223n100, 245–8, 257, 299, 329
 Quine's holism ix, 25–6, 40–1, 48, 50, 55, 59, 100–1, 112, 114, 117, 122–6
 Quine's model of knowledge ix, 39–61, 71–2, 86, 99–127
 Quine's naturalism 14, 70, 113, 122–5, 164
 Quine's neglect of reason, *see* reason, Quine's neglect of
 Quine's (radical) empiricism ix, 59, 64n, 100–3, 105–13, 124–5, 327, 342

 Railton, P. 73, 90
 rational intuition, *see* intuition, rational
 rationalism 18, 50, 157, 195, 215
 rationality 3, 8–11, 15–16, 25, 52, 58–9, 65, 152
 Rawls, J. 31–2
 realism 50, 56, 60, 73–4, 82–3, 91–100, 104, 108–9, 162, 181, 203, 206–7, 215n88, 256, 258, 261, 341
 basic, *see* basic realism
 epistemic 96–7
 existence requirement 91–2, 97–8
 flexible 92–5, 97–8
 independence requirement, *see* independence
 internal 97–8
 literalness requirement 92
 metaphysical 98, 166, 213
 rigid, non-rigid 91–5, 97–8
 robust 97
 strong 95–8
 syntax-ontology parity 92–3

 rear of knowledge, *see* knowledge, rear of
 reason 5, 6n5, 50, 73–4, 84–5, 87, 100–2, 112, 115–16, 149, 176, 249; *see also* intellect
 Quine's neglect of *x*, 100–1
 receptivity 8
 reductionism 43, 46, 102–3, 133n8, 157, 165, 230
 reference 53, 93, 133, 141, 146, 156, 159, 164, 187, 190, 200–1, 207, 216–17, 229–30
 composite 201
 simple 201
 reflective equilibrium 31–2
 Reid, T. 34
 relational holism (RH), *see* holism, relational
 relativism 109, 147, 179–80
 representation 14, 66, 139, 143n28, 170, 183–5, 191n35, 199, 202–3, 213, 278n44, 297n69
 Rescher, N. 22, 33
 Resnik, M. D. 24, 153, 208n77, 216n88, 258, 324–5
 Restall, G. 334–6
 revisability 45–7, 55, 59, 71, 74n2, 114–15, 126, 160, 331
 revision 20, 31–2, 41, 45–7, 51, 58n26, 60, 71, 74n2, 106–7, 115, 118, 122, 124, 126, 147, 236, 257, 284, 295, 327, 328–31
 in logic 46, 51, 60, 106–7, 236, 242, 257, 284, 295, 327–31
 revolution 41, 50, 54, 60, 106, 120, 123–6, 136, 180
 Copernican 15, 47, 64, 86, 167, 242
 scientific 69, 71, 204n67, 180, 340
 RH, *see* holism, relational
 Robinson, A. 322n110
 Romanos, G. D. 101
 Rorty, R. 21, 72, 98, 182–3
 Rosen, G. 35n20, 77–8, 95–96, 213n84
 Rosenkoetter, T. 136n17
 route of correspondence, *see* correspondence, route of
 Russell, B. 30–1, 85, 109n10, 140n24, 142, 151, 197n50, 208n77, 215n88, 221, 241, 246, 256, 259, 261, 279, 293, 298n70
 Russell, G. 113, 116n20
 Russell's paradox 30–1, 85, 151, 261, 293, 318n102, 323, 327

 Sagi, G. 302
 same logical operation 305, 311
 satisfaction 45n7, 146, 156, 217, 224–35, 257, 290, 300
 Schlick, M. 185n28
 science viii, *x*–xi, 5–6, 8, 10–12, 14–15, 24–6, 29, 31–2, 34–5, 39–40, 42–3, 45–6, 53–4, 56, 59, 64–7, 69–72, 74n2, 76–7, 82, 84, 86, 90–1, 93, 100–4, 106, 108–11, 117–26, 133–4, 136, 143–8, 151–2, 154,

- 157–8, 164, 167–8, 180, 194, 200, 204,
 211–12, 221, 232, 239–40, 242, 246–7,
 249, 251, 254, 258–9, 261, 292–3, 318,
 330, 340
- scientific revolution, *see* revolution, scientific
- Sellars, W. 7–8, 20, 251
- semantic ascent 63
- semantic concept 220, 257
- semantic formalism 321
- semantics 51n18, 68n9, 104, 147, 195n46,
 220–1, 228, 257, 269n31, 279–81, 284,
 299, 318, 322, 333
- Tarskian 68n9, 220–1, 228, 257, 297n69
- Sen, A. 15–16
- sensory perception, *see* perception
- set theory 31, 94, 108, 196, 198–200, 212, 223,
 226, 232, 275, 290, 298n70, 304–10, 314,
 318–24
- Shapiro, S. vii1n1, 21, 24, 207–8, 258, 324–5
- Sheffer, H. M. 243
- Sher, G. 20, 59n28, 61, 68n9, 82n11, 133,
 156n51, 160, 166n9, 227, 234n118,
 240n3, 258, 260n8, 263n12, 281,
 286n55, 297n69, 302n74, 304–5,
 307, 311, 313n91, 318–19, 321n107,
 331n116
- Tarski-Sher thesis 304–7, 311
- Sider, T. 35n20, 208n77, 259
- Siewert, C. 164
- Simchen, O. 164n4
- skepticism 4, 15, 20, 50, 54n24, 70, 138, 147,
 174–81, 184, 203, 244, 297, 341
- global 175–6, 178
- local 175–6
- moral, *see* moral skepticism
- Soames, S. 220n95
- Sober, E. 96
- Sorensen, R. 93n35
- Sosa, E. 9n8, 18n1, 23, 27–8, 30, 34
- Spelke, E. 256
- spontaneity 7, 10, 14
- standard first-order thesis, *see* first-order, thesis
- standard of truth x–xiii, 9, 30, 39, 57, 94, 96,
 147, 155, 162–3, 165, 168–79, 181–3,
 188–91, 198, 216, 218, 235, 259,
 321n108, 335–6
- Stanford, K. 121n30
- Stanley, J. 85n19
- Stavi, J. 300
- Steiner, M. 204n69
- Strawson, P. F. 46n13, 55n24, 113–16
- strict ordering requirement, *see*
 foundationalism, strict ordering
 requirement
- Stroud, B. 113, 176
- structuralism, mathematical 24, 321, 324–5
- substantive xi, xiii, 8, 11, 17–18, 21, 30, 35,
 48–9, 72, 90–1, 98, 107, 117n20, 131–4,
 136–9, 141, 143, 148, 151–2, 154,
 157–61, 174, 176, 189, 203, 219, 230,
 241–2, 260, 296, 300, 306, 327, 331,
 336, 340
- theorizing viii, xi, 133–4, 181, 315n97
- theory of knowledge 72, 137–8
- theory of truth 127, 131–43, 148, 151–4,
 160–2
- substantiveness xii, 11–12, 16, 29–30, 35,
 39–40, 72, 113, 127, 131, 147–8, 159–61,
 219, 230, 306, 340–1
- substantivism vii, x, 106n5, 127, 131–4, 148,
 153, 159, 174, 203, 211, 336, 340–1
- Suppes, P. 143n29, 145
- syntax 31, 72, 81–2, 86, 92, 167, 212, 223, 228,
 279, 281, 284, 320
- system of knowledge viii–xiii, 3, 9–10, 12–13,
 16–30, 35, 39, 48, 55–62, 70, 72, 74, 80,
 91–2, 96, 98, 102, 104, 118, 122, 125–6,
 134, 149, 162–3, 169, 193, 195, 204, 207,
 236, 240, 244–9, 253–5, 269, 302, 307,
 309, 329, 331, 333–4, 336
- systematicity 8, 11–12, 29, 62, 131, 147,
 150, 298
- T-sentence 63, 158
- Tappolet, C. 156, 216, 249n14
- Tarski, A. 32–3, 52n19, 63, 145–7, 151–3, 158,
 166–7, 178, 183, 187, 204n67, 217,
 219–34, 253–4, 256–7, 269n31, 278, 286,
 296–8, 301–2, 305, 311, 314n95, 318–19,
 321, 329–30
- Tarski-Sher thesis, *see* Sher, Tarski-Sher
 thesis
- Tarski's definition of logical
 consequence 225, 227, 257, 281, 296
- Tarski's problem 294, 296–7
- Tarski's theory of truth 32–3, 146–7, 152–3,
 158, 167, 217–34, 281, 300–1
- Tarskian hierarchy 33, 166n9, 231–4
- Tarskian semantics, *see* semantics, Tarskian
- Thagard, P. 69n11
- Thalos, M. 205, 316n99
- theorizing
- frictionless, *see* frictionless, theorizing
- substantive, *see* substantive, theorizing
- theory
- of knowledge vii, 3n1, 9n8, 17–21, 25, 43,
 46–8, 50, 55, 59, 64, 67, 71–2, 76–7, 80,
 90n28, 93, 103, 105, 124, 131, 137–8,
 148, 152, 162, 295, 340
- logical 5, 12, 35, 64, 105, 235, 240, 242,
 245–7, 251, 260–1, 264–6, 270, 293n66,
 308, 317n100, 327–30, 335
- of mathematical truth, *see* new theory of
 mathematical truth
- meta-, *see* metatheory
- philosophical 17, 41, 102, 127, 131, 138, 148

- theory of truth (conception of truth, approach to truth)
 coherence 153–7, 159, 169, 187, 189, 216
 correspondence, *see* correspondence
 deflationist, *see* deflationism; deflationist functionalist 154
 substantive, *see* substantive, theory of truth
 substantivist vii, x, 127, 131–4, 153, 159, 211, 340
 Tarski's, *see* Tarski, Tarski's theory of truth
 thing in itself 4, 98, 166, 172, 181, 183, 259, 342
 Tolley, C. 14n11, 241n7
 tonk 248, 258
 topic neutrality xiii, 236, 254, 288, 291, 294
 transcendence xi, 72–3, 135, 137, 157, 161, 164–71, 178–89, 202, 204, 213, 232–3, 258–9
 HG- 166–7
 HH- 166–7
 immanent 72, 168n11
 transcendent mode of thought xi, 166, 168, 183
 transcendental, transcendentalism 6, 8–9, 11, 14, 135, 137, 157, 164, 167
 truth
 basic principles of 162–236
 correspondence principle of, *see* correspondence
 criterion of, *see* criterion, of truth
 disunity of xi, 127, 134, 144, 153, 166, 216
 diversity of xi, 134, 143, 146–8, 160–1, 166, 189, 218, 230, 337
 epistemic approach to 162
 formal 212, 295–6
 fundamental principle of, *see* fundamental principle of truth
 logical, *see* logical truth
 logicity principle of 218–19, 229–30, 232, 234, 292, 300
 mathematical, *see* mathematical truth
 minimalist approach to, *see* minimalism
 in a model 225, 280–1, 284
 platitudes of, *see* platitudes
 question of, *see* question of truth
 schema 63
 standard of, *see* standard of truth
 structure of 127, 129–236
 substantive principles of xi, 154, 160
 substantive theory of, *see* substantive, theory of truth
 substantivist approach to 148
 Tarski's theory of, *see* Tarski's theory of truth
 predicate 32–3, 63–4, 105, 132, 160, 174, 222–3, 231–2, 257
 unity of xi, 127, 134, 157, 191, 216
 truth-apt 158, 165, 191
 truth-bearer 134n14, 136, 140, 152, 162, 165
 truth condition 96, 134, 138–41, 143, 146–7, 152–3, 156, 158, 174, 187, 202, 218, 226, 230, 234, 258, 264–5, 277, 280, 300–1, 335
 “Two Dogmas of Empiricism” ix, 40, 46–51, 101–2, 104, 113
 type
 ontological 81, 92, 315, 323
 syntactical 92, 226n106
 understanding 4–6, 9–10, 14, 84, 110, 115–16, 148–9, 158, 277
 unfixed 68, 115, 324
 unification 144, 151–2, 157–9
 unity xi, 8, 11–12, 16, 39, 65, 71, 125n32, 127, 134, 147, 149–50, 152, 154, 157–61, 191, 216, 218, 229, 249, 286, 311–12, 331, 337
 challenge 134, 157
 universals 78n8, 82, 108
 Väänänen, J. 308
 vagueness 245
 Van Atten, M. 320
 Van Benthem, J. 243, 299
 Van Cleve, J. 136n17, 183–4, 329
 Van Dalen, D. 320
 Vanzo, A. 136n17
 Varzi, A. 93n36
 veridicality vii, ix, 5, 8–9, 11–12, 17, 23–4, 28, 35, 39–40, 72–3, 100–1, 125–7, 246, 331, 340
 of logic ix, xiii, 5, 12, 235, 239–43, 245, 250–1, 259, 267, 270–1, 283, 327, 331–2, 337–8
 verificationism 48
 Wang, H. 109
 Watson, J. 15, 86
 web of belief 34, 126
 Westerståhl, D. 299–300, 313n91
 White, M. 257, 329
 Whitehead, A. N. 208n77
 Wiles, A. 86
 Williams, B. 172, 295
 Williams, M. 22
 Williamson, T. 73, 88n24, 90
 Wittgenstein, L. 4, 6–12, 21, 53, 72, 100, 141, 150–1, 164, 243, 256, 342
 Woodward, J. 288
 Wright, C. 73, 133, 153–4, 156–60, 170, 187, 243, 320
 Wright, C. D. 132–3, 153–4
 Yablo, S. 164n4, 213
 Yi, B. 202n62
 ZFC 199, 275, 279, 284, 286–7, 298n70, 302n74, 313–19, 322