

THE
CONSTITUTIVE
A PRIORI

ARTHUR SULLIVAN

DEVELOPING AND EXTENDING AN
EPISTEMOLOGICAL FRAMEWORK

The Constitutive A Priori

The Constitutive A Priori

Developing and Extending an Epistemological Framework

Arthur Sullivan

LEXINGTON BOOKS

Lanham • Boulder • New York • London

Published by Lexington Books
An imprint of The Rowman & Littlefield Publishing Group, Inc.
4501 Forbes Boulevard, Suite 200, Lanham, Maryland 20706
www.rowman.com

Unit A, Whitacre Mews, 26-34 Stannary Street, London SE11 4AB

Copyright © 2018 by The Rowman & Littlefield Publishing Group, Inc.

All rights reserved. No part of this book may be reproduced in any form or by any electronic or mechanical means, including information storage and retrieval systems, without written permission from the publisher, except by a reviewer who may quote passages in a review.

British Library Cataloguing in Publication Information Available

Library of Congress Cataloging-in-Publication Data

Names: Sullivan, Arthur, 1970- author.


Title: The constitutive a priori : developing and extending an epistemological framework / Arthur Sullivan.

Description: Lanham : Lexington Books, 2018. | Includes bibliographical references and index.

Identifiers: LCCN 2018006089 (print) | LCCN 2018002345 (ebook) | ISBN 9781498547123 (electronic) | ISBN 9781498547116 (cloth : alk. paper)

Subjects: LCSH: A priori. | Knowledge, Theory of.

Classification: LCC BD181.3 (print) | LCC BD181.3 .S85 2018 (ebook) | DDC 121/.4—dc23

™ The paper used in this publication meets the minimum requirements of American National Standard for Information Sciences—Permanence of Paper for Printed Library Materials, ANSI/NISO Z39.48-1992.

Printed in the United States of America

Contents

Preface	vii
Acknowledgements	xi
PART I: INTRODUCTION	
1 A Sketch of the Terrain	3
2 Further Preliminaries	25
PART II: KEY VARIETIES OF IMMUNITY TO COUNTEREXAMPLE	
3 Necessity and Analyticity	49
4 <i>A Priori</i> Justification	75
PART III: ADAPTING AND APPLYING THE CONSTITUTIVE <i>A PRIORI</i> APPROACH	
5 Two Major Challenges to the <i>A Priori</i>	99
6 Modal Revisionism and Moderate Externalism	123
PART IV: MAPPING OUT A CONSTITUTIVE <i>A PRIORI</i> VIEW	
7 Entailments and Conclusion	161
Bibliography	177
Index	185
About the Author	189

Preface

The question of the *a priori* is a core issue running throughout the history of philosophy. Variations on this theme include: Is there a non-empirical source of justification? Can an adequate epistemology be developed without appeal to one? What is the precise nature, and status, of non-empirical justification? As the cases of at least Plato, Descartes, Kant, and Frege illustrate, many seminal original philosophical positions have grown out of work on questions surrounding *a priori*. Further, on some (though by no means all¹) conceptions of the discipline, *a priori* is the very essence of philosophy, as distinct from other theoretical enterprises—and so skepticism about the *a priori* is tantamount to pessimism as to whether philosophy has a distinctive subject matter or methodology. Philosophers' views about *a priori* are inextricably linked not only to various other issues in epistemology, semantics, and metaphysics, but also to their conceptions of the discipline of philosophy as a whole—including in particular its proper methodology, and its scope or range.²

Recent decades have seen some provocative and potentially epochal work on *a priori*, which has served to both refine and to complicate maps of the terrain, including: (i) Kripke's (1972) conjectures about contingent *a priori* and necessary *a posteriori* statements, and the flood of discussion they sparked;³ (ii) the range of critical evaluations of Quine's (1951) influential but no longer sacrosanct attack on the analytic/synthetic distinction;⁴ and (iii) developments in the semantics of indexicality, and in multi-dimensional semantics.⁵ One primary aim of this book is to make some contributions toward a clear-headed evaluation of the upshot of these developments for the issue of *a priori*.

I will argue that the notion of the constitutive *a priori* provides the best means, all things considered, of accommodating these recent developments into a coherent, compelling view. The constitutive *a priori* is a broadly neo-Kantian orientation on the *a priori*, versions of which were pioneered by

Reichenbach (1920), Wittgenstein (1921, 1953, 1969), Carnap (1937, 1950), and Pap (1946). Some other terms in the literature for what I will call the ‘constitutive *a priori*’ include the ‘relativized’, ‘functional’, ‘dynamical’, and ‘adaptable’ *a priori*; and C.S. Pierce, Henri Poincare, Ernst Cassirer, and C.I. Lewis are among those who receive mention as seminal influences on the tradition. Recent work within this orientation which significantly influences this present project includes Friedman (1992, 2000, 2007, 2011), Railton (2000, 2003), and Stump (2003, 2011, 2015).⁶

The constitutive *a priori* is most commonly known as a position within the philosophy of science, the guiding idea being that one of Kant’s (1781) signature moves provides the means to incorporate unforeseen drastic shocks into existing theory. (The two key catalytic drastic shocks, within the tradition, are non-Euclidean geometry and the theory of relativity—cf. Stump (2015) for a thorough account of the development of the constitutive *a priori* within the philosophy of science.) I aim to show that the notion of the constitutive *a priori* provides not merely a satisfactory epistemological framework, but, further, a compelling way to accommodate and integrate some of the most significant lessons learned in twentieth-century philosophy. The distinctive original contribution of this present work lies in the case it builds for taking this constitutive *a priori* orientation as a good means of integrating and consolidating certain epochal insights of Wittgenstein, Carnap, Quine, Kripke, and Kaplan.

Down a confluent avenue, Soames (2003: Vol. I, xi) asserts that one of the most significant achievements of twentieth-century philosophy is precisely the ‘success achieved in understanding, and separating from one another, the fundamental ... notions of logical consequence, logical truth, necessary truth, and *a priori* truth.’ One of the strands within this general trend which I will develop in this book, and which plays a role in defining its structure, is a judicious distinction between the *a priori* and two closely related notions—that is, the metaphysical notion of necessity and the semantic notion of analyticity. Developing and defending a suitably refined notion of *a priority* depends on a clear-headed and refined view of the relations between properly epistemological considerations and distinct questions and issues in bordering parts of semantics and metaphysics.

[§]

Here is the plan. Part I gives a working, reasonably comprehensive introduction to the terrain. The question of *a priority* is developed from a conceptual point of view (in §1.1) and from a historical point of view (in §1.2). The rest of chapter 1 and all of chapter 2 are dedicated to fostering technical precision when it comes to certain distinctions and refinements, pertaining to concepts, thought, and language.

Part II then consists of separate analyses of the substance of the metaphysical concept of necessity, the semantic concept of analyticity, and the

epistemic concept of *a priori*, and considers arguments for and against the intelligibility and worth of each concept.

Part III is an exploration of what I take to be the two main recent shocks to the *a priori*—that is, the challenge of revisability (associated with Quine, circa 1950) and the externalist challenge (associated with Kripke, circa 1970). Here the original philosophical contributions of this present research begin to really develop, in the course of working through these epochal challenges. It is here that the contours, and the benefits, of a constitutive *a priori* view get drawn out in an innovative way.

Finally, the aim of Part IV is to map out a constitutive *a priori* view. I address the precise entailment relations between the concepts of necessity, analyticity, and *a priori*. I also consider general questions about the relations between the three subfields of philosophical inquiry in which they belong, as well as some of the perennial philosophical questions of which these three concepts are core ingredients.

[§]

A final note about the structure of the book: Parts I and II might appear to be mere stage-setting before getting into the heart of the matter; if so, why are they so long? Why does the stage-setting take up more space than the featured play? My answer as to why so much careful attention is paid to conceptual clarity and fine distinctions may well be typical of philosophy generally—or, at least, of the kind of philosophy which takes twentieth-century philosophy of language to have constituted an important step in the maturation of this ancient discipline. These preliminaries are essential to defining concepts and propositions which are sufficiently precise as to promise worthwhile engagement. In particular, it will take a lot of careful digging to defend *a priori* (and some of its fellows, such as analyticity) from some of the considerable mass of problems and objections which have been raised against it.⁷

By the end of Parts I and II, I will have accrued a precise toolkit, an apparatus of instruments, which will then be available for further jobs. In this way the explanation for the excessive length of Parts I and II also and thereby accounts for the relative brevity of Parts III and IV; those latter tasks can only be dealt with after the painstaking work done in Parts I and II. The glamorous work of developing the original philosophical contributions of this present research can only happen once the relatively dry and dusty conceptual housecleaning is properly executed.

NOTES

1. Williamson (2007, 2014), for example, opposes this conception of philosophy; though it is noteworthy that he clearly takes himself to be swimming against the prevailing tides.

2. Cf. Peacocke (2006).
3. Cf. Soames (2003: Vol.II, Part 7) for a survey of these issues.
4. Cf. Boghossian (1997), BonJour (1998: Ch.3), Sober & Hylton (2000), Casullo (2003: Ch.5), and Russell (2008: Part II) for a variety of good critical discussions of exactly what Quine's (1951) arguments succeed in establishing.
5. Here cf. especially Kaplan (1989), Stalnaker (2001), Chalmers (2006).
6. Some other recent work within this tradition includes DiPierris (1992), DiSalle (2002), Richardson (2002), Franco (2011), Angeloni (2012), Uebel (2012), and Everett (2015).
7. Analyticity is a case in point, in this respect. Only given the careful work on semantic minutia, particularly in §§s 2.1, 2.2, and 3.2, can the case against analyticity be considered and addressed in §3.3.

Acknowledgements

I have worked on this book, on and off, for over two decades, and so have received much and various help along the way. Adèle Mercier and Henry Laycock were both great influences on my philosophical world in graduate school, and will find some of themselves in here. Conversations with James Cargile and Nenad Miscevic, very early on in the project, have proven to be more helpful than they could guess. Over the past decade I have taught three different seminar courses on ‘Twentieth-Century Developments on the A Priori’ here at Memorial University, each of which has resulted in significant improvements to the depth and breadth of my research; I am grateful to my students for their engagement and their efforts.

Thanks to the Social Sciences and Humanities Research Council of Canada for support—some portion of two Standard Research Grants has funded travel to present research, as well as graduate student help with the project. (Huge thanks to Michael Lahey, who was the biggest help with the latest stages of the project.) Material in this book was presented at (at least) Logic & the Philosophy of Science III at Bogata in 2014, and at three different meetings of the Canadian Philosophical Association (at Ryerson University in 2017, the University of Victoria in 2013, and Concordia University in 2010). Thanks to commentators and audiences for their questions and challenges.

Finally, thanks to my editor Jana Hodges-Kluck and the rest of the staff at Lexington Books/Rowman & Littlefield, as well as to Jayanthi Chander and the production team at Deanta Global, all of whom have proven to be exceedingly conscientious, in the process.

Part I

INTRODUCTION

Chapter 1

A Sketch of the Terrain

§1.1: THE QUESTION OF A PRIORITY

'*A priori*' (from the Latin word which gives us *prior*, or before) contrasts with '*a posteriori*' (*posterior*, or after); the terms are used to mark a central fissure in epistemology. In contemporary discussion, '*a priori*' is used to designate a putative non-empirical source of justification.¹ Debates surrounding the existence and nature of the *a priori* thread through the history of philosophy, connecting with and underlying various other questions, issues, and debates. Broadly speaking, the many varieties of rationalists believe in some or other version of the *a priori*, as do many moderate empiricists. In contrast, the very idea of non-empirical justification is dismissed as, at worst, incoherent, or, at best, unhelpful, by many proponents of radical empiricism, naturalism, and skepticism.

At first pass, something is justified *a priori* if and only if the grounds for believing it do not depend on sensory evidence. A plausible candidate for being justified *a priori* is:

1. One cannot steal one's own property.

To determine whether one is justified in believing that [1] is true, one need not take a poll to find out how many people agree, or devise a variety of experiments to test whether one could succeed in stealing one's own property.² One's justification for believing [1], it seems, has nothing to do with that kind of evidence. Some of our beliefs are justified empirically (e.g., 'There is an apple pie baking nearby,' based on an olfactory experience; 'This piano is out of tune,' based on an auditory experience). Most philosophers hold that not all of our knowledge can be so justified, and so posit the

category of *a priori* justification. In particular, from at least Plato on down, many have held that some of our knowledge is universal, necessary, and simply immune to counterexample, and that such knowledge could not possibly be justified empirically.

Truths of logic (such as [2]), mathematics ([3]), and ethical principles ([4]) are commonly cited candidates for instances of *a priori* justification:

2. No proposition is at once both true and not true.
3. Two is a factor of every even number.
4. One ought to keep one's promises.

We will consider these and other candidates at greater length in due course.

[§]

The epistemological notion of *a priori* is a close relative of the metaphysical notion of necessity and the semantic notion of analyticity. One central thing that these three concepts have in common—and which, to a large extent, accounts for why they are of deep and enduring philosophical interest—is that all are tightly linked, in some way or other, to what we might call *immunity to counterexample*.³ Further, little attention was paid to such fine distinctions between these three concepts, traditionally, because it was largely presumed that they are co-extensive.⁴ So, for example, contemporary philosophers may be inclined to read Hume's (1748: Sect. IV, Part 1) discussion of the category of 'Relations of Ideas' as involving a murky, indiscriminate mixture of epistemological, semantic, and metaphysical considerations. (Much the same could be said of Hobbes' or Locke's discussions of 'trifling propositions'⁵—even more so of Plato's epochal discussions of the Forms.) For another example, Pap (1958: Ch.1) documents ways in which Leibniz uses interchangeably terms which no longer seem so obviously tightly linked—including 'necessary,' 'universal,' 'eternal,' and 'certain.' Whether it be Descartes' (1641) guiding assumption that 'conceivable' entails 'possible,' or Kant's (1781: B15) presumption that necessity is a criterion for *a priori*, variations abound on the theme that *a priori*, analytic, and necessary are but different aspects of the same phenomenon.⁶

In recent decades, there have been some grand shocks to traditional presumptions about the relations between these concepts. One blow comes around 1950 with attacks by Quine (among others, but most famously) on the analytic/synthetic distinction.⁷ Many have felt that Quine's arguments against analyticity are decisive, and also raise serious concerns about the concepts of *a priori* and necessity; others have responded with serious questions both about the cogency of Quine's argument against analyticity, and about the allegation that these arguments even remotely touch on those

other concepts. Another jolt comes around 1970, with Kripke's arguments, in the wake of some groundbreaking arguments for semantic externalism, that necessity and *a priori* are not co-extensive concepts—that is, that there are necessities which are not knowable *a priori* and things knowable *a priori* that are not necessarily so.⁸ Kripke poses a challenge not only to certain kinds of rationalism, which hold that one can infer necessity from *a priori* or vice versa, but also, more generally, to a wide variety of inferences which immediately, or implicitly, slide back or forth among such roughly overlapping but distinct concepts. Kripke's arguments, too, have provoked a flood of discussion. A third, related shock to the traditional world-order concerns developments in the semantics of indexicality, and in multi-dimensional logics—see especially Kaplan (1989), Stalnaker (2001). This work has further complicated relations between various different senses in which something may (appear to) be immune to counterexample. One goal of this book is to make a contribution toward a clear-headed evaluation of the coherence and significance of *a priori*, and its relations to necessity and analyticity, in the wake of these developments.

[§]

For starters, I will begin the job of sketching fairly precise working conceptions of the terms 'necessary' and 'analytic.' The goal is to specify concepts that are refined enough so that specific theses about their interrelations with *a priori* can be tested, but yet are not out of touch with their broad historical roots. I will start this job in a preliminary way here; this part of the ongoing project will continue through Parts I and II.

Necessity is a concept of metaphysics. To say that something is necessary is to say that it could not fail to be—regardless of time or place, and irrespective of how contingent matters of fact may be altered. Necessary truths are absolutely unalterable, and we have no choice or influence in the matter. Putative examples of necessary truths include:

5. No two solid objects can simultaneously occupy the same spatial location.
6. Increasing an object's velocity will increase its momentum.

As Kripke (1972) stresses, necessity, *per se*, has nothing to do with anyone's knowledge of anything. That something is necessary does not immediately entail any epistemological or semantic conclusions. Necessary truths are a matter of how things stand in mind-independent reality, and are something about which all humans could be mistaken or ignorant.⁹

As distinct from both '*a priori*' and 'necessary,' 'analytic' is a semantic notion. First and foremost, it has to do with meaning, or semantic content. A truth is analytic if and only if it is true by virtue of the meanings of the

terms involved; alternatively, the denial of an analytic truth is contradictory. Candidates for analyticity include:

7. Squares have four equal sides.
8. No grandmothers are childless.

To call something analytically true is to say that the meanings of its parts, and the way in which they are combined, suffice to ensure that it is immune to counterexample. So, it is not hard to see why the notion of analyticity has been put to work in both epistemology and metaphysics. (After all, there are close constitutive connections between linguistic meaning and the content of our beliefs, as well as between meaning and the language-independent objects of most of our thought and talk.) On the epistemic front, provided that just grasp of the meanings involved is sufficient for recognizing that an analytic statement is true, it is plausible to think that analyticity may hold the key to a satisfactory theory of the *a priori*. On the metaphysical front, there have always been philosophers who are wary of the idea that necessity is an objective feature of mind-independent reality. Down this avenue, there is a temptation to think that perhaps the source of our intuitions of necessity resides in language, and in our conventional categories—that is, that what many have mistakenly thought to be metaphysical necessity is more clearly and helpfully characterized as analytic truth.

[§]

To sum up, let me stress the *prima facie* difference between these concepts. First, to ask whether something is necessary is to ask: Could this possibly be false, should contingent matters go (or had gone) otherwise? Second, to ask whether something is analytic is to ask: Does the meaning of the constituent bits, plus the mode of composition, suffice to ensure that this is true? Lastly, to ask whether something is *a priori* is to ask: Are we justified non-empirically in believing this to be true?

There will obviously be considerable overlap between these three concepts; but the traditional presumption that they are but different guises of the same underlying phenomenon is no longer very plausible. In Part IV, to help develop my own favored constitutive *a priori* approach to the question of *a priority*, I will chart interrelations between these three concepts.

§1.2: SOME HISTORICAL BACKGROUND

The following is an outline of the broader history of some relevant philosophical debates. It sketches some core themes running through much of

Western philosophy, in a way that will be used to structure much of the succeeding discussion.

Necessity is a metaphysical concept. To say that something is necessary is to say that it could not fail to be—regardless of time or place, and irrespective of how contingent matters of fact may be altered. Necessary truths are absolutely unalterable, and we have no choice or influence in the matter. The concept of necessity is integral to systematic inquiry. Many great leaps forward for human knowledge take the form of discovering that some phenomena—such as fertile soil, lightening, or the plague—are not in fact randomly scattered and unconnected from the rest of nature, but are rather intrinsically related to other phenomena by general principles that are rooted in the very nature of things. In one crucial sense of ‘understanding,’ we understand some phenomenon when we know where it stands in a nexus of causes, laws, and effects such that, given the antecedent conditions and the laws of nature, the phenomenon is necessitated: it could not but happen. Necessary truths, those which capture the features of the world that could not be otherwise, are a very special subset of the set of truths, a set in which philosophers—as well as mathematicians, scientists, and many others, of course—have long been interested.

The epistemological notion of *a priori* (i.e., justification which does not depend on experiential evidence) is also central to systematic thought about human knowledge. Indeed, this notion of that which can be known to be true just by thinking about it, as opposed to requiring some process of gathering evidence, is sometimes taken to be definitive of the subject matter of the discipline of philosophy. A prevalent (though by no means unanimous or uncontroversial) idea is that discovering or justifying things via thought is to philosophy what discovering or justifying things via the experimental method is to science.

An ancient theme in philosophy concerns the fit that exists between the two concepts of necessity and *a priori*. In one direction, there is good reason to think that only necessary truths can be known *a priori*: that is, it is hard to see how one could obtain knowledge of accidental contingencies, of matters of fact that could very well have been otherwise, just by thinking through the content of our concepts. In the other direction, it also seems fairly evident that the only things which can be known *a priori* must be necessary truths—for if something is contingent, and so varies according to circumstances, how could we know that it is true just by reflecting on the content of our concepts?

In due course we will see that both of these putative connections have been questioned (i.e., there are plenty of reasons to doubt whether all, or only, metaphysical necessities are knowable *a priori*). However, our immediate concern is with a related but distinct thesis about the connection between necessary and *a priori*, which has long been a matter of controversy: Plato believed, as many have since, that we have *a priori* knowledge of necessary

truths, that our faculty of reason permits us direct access to certain necessary features of reality. Since many pertinent debates within epistemology and metaphysics date back to Plato's thought, it will prove worthwhile to explore some of his ideas.

[§]

One key problem that leads Plato to make this claim (i.e., that we have *a priori* knowledge of necessary truths) is to explain how it is that some of the things that we know seem to be fashioned from the hardest of steel—that is, to be timeless, self-evident, and, quite simply, not conceivably false. I will call this 'Plato's problem':

Plato's problem: Given that our personal experience is confined to a particular limited perspective, how can we account for our evident knowledge of timeless general certainties?

Take the simplest of examples, '2+2=4,' to illustrate. Such truths exhibit a remarkable indifference to contingent facts and to psychological processes: the facts could change—about how humans count, or how computers calculate, for example—but neither of these sorts of concerns would have any effect on whether two of anything added to two other things equal four things.¹⁰ Before anyone ever thought of it, two of anything added to two of anything yielded four; even if most people were mistaken about it, that mistake would not take one of the four away, or add one to it. What makes it true that two of anything added to two of anything yields four does not depend on any particular fact, or on how many people agree with it. What makes it true is, it seems, a general necessary feature of the mind-independent nature of things.

Plato—among many other subsequent thinkers—held that many other kinds of truth, not just mathematics or logic, also have this status of being necessarily true and knowable *a priori*. (For instance, many have thought that moral truths, such as that one ought to keep one's promises, or that murder is wrong, are *a priori* necessities.) The important point for now is the problem raised by our knowledge of these certainties, which are fashioned from the hardest of steel. Some of the things that we know, it seems, are entirely indifferent to particular matters of fact, and are not determined or constrained by the particular perspective from which the knower arrives at them. Rather, they just cannot be false. Plato's problem is to explain the nature and ground of these hardest-of-steel truths, and of our grasp of them, given the particular and limited nature of our experience. All human experience is finite, particular, contingent; yet some human knowledge is universal and necessary. What must our minds be like, in order that we can come to know such things? What must mind-independent reality be like, in order that such things are knowable about it?

Plato's seminal answer consists of the metaphysical claim that mind-independent reality must be underlain by an unchanging necessary order, and the epistemological claim that the human mind must have some channel of direct access to this order. His view is that we can attain rational insight into the timeless general features of reality, in virtue of a certain distinctive cognitive faculty, which is capable of *a priori* knowledge of necessary truths. Here is an excerpt from the dialogue 'the *Phaedo*' in which Plato discusses this faculty, in the course of exploring our ability to grasp such concepts as justice, beauty, and goodness:

[One] attains to the purest knowledge of them who goes to each with the mind alone, not introducing or intruding in the act of ... any other sense together with reason, but with the very light of the mind in her own clearness searches into the very truth of each.

[To attain such knowledge one must have] got rid, as far as [one] can, of eyes and ears ... these being ... distracting elements which when they infect the soul hinder her from acquiring [this kind of transcendent, non-sensory] knowledge.

And thus having got rid of the [limitation to sensory evidence] we shall be pure to hold converse with the pure, and know of ourselves the clear light everywhere, which is no other than the light of truth. (65e–66b)

In the first paragraph here, Plato contrasts this rational faculty of mind, this 'very light of the mind in her own clearness,' with the senses (i.e., empirical justification). (It is common among subsequent rationalists to follow Plato in speaking of rational intuition as a kind of mental seeing.) In the second paragraph he asserts that one must ignore the information coming through sensory channels, one must turn off these distractions as much as is possible, in order to really get this rational faculty of mind up and running. In the final paragraph he says that, insofar as one is able to do so, one is then free to mentally tap into the realm of eternal necessary truth.

Let us distinguish these two constituent elements of Plato's solution to Plato's problem:

Plato's metaphysical claim: there are timeless unchanging features of mind-independent reality, underlying the fleeting appearances which we experience, and grounding necessary truths. (This claim is known as 'metaphysical realism.')

Plato's epistemological claim: our minds are endowed with a special faculty, distinct from the senses, by means of which we have cognitive access to necessary truths. (This non-sensory faculty of mind is commonly called 'rational intuition,' and commitment to it is called 'rationalism.')

Note that both terms ‘metaphysical realism’ and ‘rationalism’ are used rather diversely, and these are but one of many technical senses of each term. To be sure, though, these are, both historically and conceptually, central senses of these terms.

For instance, Descartes (1641) and Russell (1912) are both rationalists, in that they hold that we have *a priori* access to necessary truths, but this does not imply that they accept all of Plato’s views on the precise nature and workings of rational intuition, or that they themselves do not disagree on pertinent questions. There is lots of room for substantive disagreements among rationalists, as I am using the term. In particular, even though—by definition, in this sense—all rationalists would accept Plato’s distinction between rational intuition and sensory justification, many succeeding, more moderate, rationalists would reject the extravagant claims Plato makes in the second and third paragraphs in the above excerpt. (There is much more extensive discussion of this issue below, first in §1.4 and subsequently throughout Parts II and III.)

[§]

Plato’s realism and rationalism have always had their detractors. Both elements of Plato’s platform have been judged by many to be beyond belief—that is, skepticism surrounds both the metaphysical claim that there are necessary truths lurking out there, beyond the bounds of space and time, and the epistemological claim that we have some faculty for gathering knowledge, in addition to our senses, which taps us into them. What are these queer truths, exactly, that the rationalist claims to be able to see, and what, precisely, is the faculty that provides these obscure visions? Many find the appeal to rational intuition to be hopelessly obscure and unexplanatory, and to just simply not come close to a satisfactory solution to Plato’s problem. Several different philosophical debates, over the ages, have turned on these and related disagreements between Platonists and their opponents.

Empiricism, the doctrine that all knowledge comes from experience, is a longstanding opponent of rationalism. This defining tenet of empiricism stems precisely from skepticism about the appeal to rational intuition. Empiricists’ criticisms, and empiricists’ own distinctive problems, will come up at several junctures in this book. Philosophical naturalism is a variant of empiricism that is, currently, widely associated with Quine’s (1951, 1960, 1969) influence. Naturalism involves privileging the natural sciences as the highest or best sort of knowledge that humans have ever developed, and so holds that other disciplines (including philosophy) ought to remain consistent with the natural sciences, as much as is possible. Naturalism is opposed to the positing of other-worldly Platonic objects, or mysterious occult faculties of mind, on the grounds that science gets on fine without them, and they do not seem to have aided or abetted philosophical progress. (There exists nothing

supernatural.) Naturalistic pessimism as to the coherence and worth of *a priori* will come up at several junctures below.¹¹

Some opponents of Platonism have taken a dismissive stance on Plato's problem—their response is to deny that we ever attain knowledge of timeless general certainties. One key benefit of this skeptical stance is that one avoids having to face certain epistemological and metaphysical problems; however, the price is that one cannot invoke the notions of necessary truth or *a priori* knowledge. Most philosophers hold that, despite the difficult problems they bring with them, these notions are indispensable to a satisfactory account of our knowledge of the workings of the world. There are considerable reasons to hold that immunity to counterexample just simply is a brute, undeniable datum.

So, many opponents of Platonism have nonetheless felt the pull of Plato's problem. There is a strain of anti-Platonism that concedes that we do attain absolute certainty, at least (but not necessarily only) in the case of mathematics and logic; and so, even though Plato's answers may raise more problems than they solve, still there must be an explanation of these truths, and of our grasp of them. Thus, many attempts have been made to explain how it is that some of the things we know are immune to counterexample, in a way that does not posit obscure faculties of mind intuiting queer abstract objects. Here is where the notion of analyticity gets added into the mix.

[§]

In contrast to both 'necessary' and '*a priori*,' 'analytic' is a semantic notion. Just grasp of the meanings involved is sufficient for recognizing that an analytic statement is true. For this reason, many have thought it plausible that analyticity holds the key to Plato's problem: we have *a priori* knowledge of necessary truths, so the idea goes, because those truths are analytic. The hope is that analyticity provides an explanation of *a priori* which is less obscure than rational intuition: that is, all truths knowable *a priori* are analytic, and that they are analytic explains why they cannot be false.

Hence, within this strand of anti-Platonism, the tight link between *a priori* and analyticity develops into a semantic, conventionalist solution to Plato's problem: it is the conventional regularities governing the usage of our words that determines what is necessarily true and knowable *a priori*. Analyticity holds some promise to explain how it is that some of the things we know exhibit this remarkable indifference to contingent facts and to psychological factors. Many proponents of this line of thought hold that the air of profundity surrounding Plato's problem is confused obscurantist puffery, that Plato's problematic metaphysical and epistemological claims—that is, about a mysterious realm of truths and about a mysterious kind of mental access to it—could be avoided by a better semantic theory. In short, many within the empiricist tradition have thought that analyticity could afford a

simple, comprehensive solution to certain epistemological and metaphysical questions. All that is worth wanting about Platonic rational intuition can be delivered by the humbler and less problematic faculty of understanding (i.e., mere competent grasp of meaning).

To sum this up, then: Well into the twentieth century, those who take Plato's problem seriously—that is, those who reject the relevant form of skepticism, and so concede that we do attain knowledge of some certainties—fall, roughly, and with many variations, into two main groups. First, there are the many varieties of rationalism, which explain *a priori* knowledge by positing a faculty, distinct from the senses, that acquires such knowledge. Second, there are those who seek explanations of our grasp of certainties without appeal to rational intuition (a central paradigm case of which are those who appeal to analyticity to solve Plato's problem).

[§]

There is one further important stop on this orienting historical overview, as the view I develop and defend in Parts III and IV relies on Kant's (1781) innovative approach to Plato's problem. Kant initiates a radical stance concerning the relation between mind and world. At (1787, [Bxvii]), Kant uses the following analogy to explain the hypothesis: astronomers prior to Copernicus had assumed that they are stationary and that the stars revolve around them; Copernicus tried the hypothesis that observers on earth are in motion and the stars are, relatively, at rest, and the result was a great leap forward in our understanding of the universe. In the analogy, the empiricism which Hume had followed to its skeptical end—more specifically, the notions that the mind is a passive recipient of information, and that all knowledge consists of sensory input and inductions therefrom—is compared to pre-Copernican astronomy. Kant's innovative Copernican hypothesis is the idea that the mind plays an active role in synthesizing and categorizing sensory input—that is, the mind is the director, as opposed to the audience, in conscious experience.

This Copernican Turn is precisely what puts the 'constitutive' in the 'constitutive *a priori*.' For Kant here foregrounds what has become known as the 'constitutive powers of mind'—that is, the active, structuring role which the mind can be seen to play in conscious experience. I will argue that this move, as interpreted and developed by certain figures in mid-twentieth-century philosophy, affords the grounds for the best overall solution to Plato's problem. (Here I am following Coffa [1991: 263], who reads Carnap and Wittgenstein as extending Kant's Copernican Turn from its original target of experience and epistemology on into the realm of meaning and semantics.)

Hume is taken by Kant to have demonstrated the impossibility of grounding in experiential input the consistency and regularity of the world, thereby proving that empiricism can afford no satisfactory answer to Plato's problem.

Kant's hypothesis is that, to some extent, the mind constructs and constitutes this consistency and regularity. Kant provides a theory of a non-empirical source of knowledge: there are active faculties of mind, constitutive powers of mind, which categorize and give form to the matter supplied to mind via the senses, and thereby structure and constitute our knowledge.

The cornerstone of Kant's philosophical edifice is the synthetic *a priori* judgment. According to Kant, synthetic *a priori* judgments are logically possible in that the concept of a synthetic judgment does not entail or contain the concept *a posteriori*, and neither does *a priori* entail or contain analytic. Synthetic *a priori* judgments are possible for agents like us, given the Copernican hypothesis that the mind is an active synthesizer of sensory input. Kant (1781) is dedicated to establishing the legitimacy of these synthetic *a priori* judgments, on the grounds that this hypothesis does something that none of the alternatives can: namely, account for our actual knowledge. Given the hypothesis that some of our judgments are contributions to, rather than inductions from, experience, there is conceptual space for judgments which are substantial extensions of human knowledge, unlike trivial analytic judgments, but which are not given in, or generalizations from, experience.

According to Kant, as in the case of Copernicus' Revolution, what we have here is a huge theoretical step forward. As the constitutive powers of mind ground synthetic *a priori* judgments, synthetic *a priori* judgments ground the process of inquiry itself. For the first time, we have an explanation of how pure reason—*a priori* cognition—can increase the volume of human knowledge. Prior to Kant, there is analysis, an *a priori* activity that cannot really tell us anything we did not already know (according to Kant, anyway, though as we will see in §3.2, Frege [1884] and others will depart from Kant on this point), and there is observation and induction. Hume showed us how little knowledge can be secured by those two faculties, unaided. What Kant believes he has discovered is a type of *a priori* knowledge, which is universal and necessary, but yet which has real empirical bite, which can constitute significant extensions of human knowledge. For Kant, all significant scientific advances (such as the principle of the conservation of matter) and all mathematical truths (such as ' $a^2 + b^2 = c^2$ ') are synthetic *a priori* judgments—that is, they go well beyond experiential input, and they add to our knowledge of the subject concept. All significant advances in inquiry consist in the discovery and elucidation of synthetic *a priori* judgments, judgments such as 'Every event has a cause' and 'Humans are free agents who are subject to moral law.' Given the possibility of these universal necessary judgments that constitute genuine extensions of the body of human knowledge, their importance to philosophy or to human knowledge generally cannot be overestimated.

Developing exactly how my own constitutive *a priori* view departs from the letter of Kant's theory is one of the main orders of business of Parts II and III below.

§1.3: METAPHYSICS, SEMANTICS, EPISTEMOLOGY

As has already become evident, our inquiry sprawls across many distinct but interconnected subfields of philosophy. Chief among those are epistemology, semantics, and metaphysics; so some brief orienting remarks on those subfields is in order. For one thing, there is some diversity and fluidity as to how people understand and employ these terms, and so it is important to stipulate precisely how I will use them. Further, this should help to clarify some theses about their interrelations which are subsequently developed. One of my goals in this book is to employ a certain stance on *a priority* to make a contribution towards a refined conception of the nexus at which these three core subfields within philosophy overlap—particularly with a view to incorporating the many important lessons learned in twentieth-century philosophy.

A point that pertains to the structural order of presentation: Up to here, the narrative order has gone: (i) necessity, then (ii) *a priority*, then (iii) analyticity. That order is historically accurate—after all, recognition of the importance of semantic questions per se emerges at a relatively late point in philosophical sophistication. However, this order will be henceforth changed, to put semantics in the middle, between metaphysics and epistemology—where it belongs conceptually, as opposed to historically. That is the distinctive geography of the semantic, between mind and world, as meanings are intricately connected both to our beliefs and intentions and to mind- and language-independent objects which most of those beliefs and intentions concern or involve.

[§]

Etymologically, the term 'metaphysics' relates to the book that is situated within Aristotle's corpus after the *Physics*. The *Physics* is a study of nature and of natural changes, events, and forces. In the *Metaphysics*, Aristotle pursues more general and abstract questions. He refers to his pursuit as the study of being qua being—that is, the study of the most general and necessary features something must have in order to count as a being. What sorts of thing exist? Into what categories do they fall? What properties are there in common among existent things, or among these categories?

Over the centuries, 'metaphysics' has come to be applied, more generally, to inquiries into the nature of reality that go beyond or behind the methods of science. Whereas scientists aim to characterize some specific range of objects, events, and forces, metaphysicians address such questions such as:

What *are* objects, or events, or forces? The scope of metaphysics is broader than that of the physical sciences—examples of questions that are not scientific but metaphysical include:

1. What is causation?
2. What are moral values?
3. Are numbers mind-independent, or human constructions?

Metaphysics, then, names a rather large sub-field of philosophy, which overlaps with many different strains of inquiry. Whether a certain thing or kind of thing exists, how to distinguish the essential from the accidental properties of something, whether or not distinct things belong in the same general category—these are (at least in part) metaphysical questions. Some core metaphysical debates, through the ages, concern the nature of space and time, exactly what sorts of things minds are, or whether all that exists are concrete particulars. Quite generally, though, metaphysical questions are prone to arise in the course of virtually any inquiry—be it scientific (e.g., Do electrons really exist?), medical (e.g., What exactly is a virus?), legal (e.g., Who exactly count as persons?), and so on.

Semantics is the study of meaning, of the relations between linguistic signs and what they mean or represent. Philosophers have long been interested in semantics because the abilities to represent something specific in thought, and to prompt someone else to consider what one is representing in thought by producing sounds or symbols, are fundamental and distinctive human capacities. So, insofar as we want to know what makes up our cognitive lives, and what makes possible many of the distinctive achievements of our species, some understanding of the semantic properties of thought and language is indispensable.

Semantics is one core component within the sprawling interdisciplinary inquiry into our capacities to represent and communicate using the medium of language (which also involves linguists, psychologists, literary theorists, computer scientists, and so on). Some of the greatest leaps forward in twentieth-century philosophy have come in the field of semantics, and an important legacy of these developments is the explicit appreciation that there is a semantic dimension to any philosophical issue or question.

The word ‘epistemology’ means ‘the study of knowledge.’ The aim of any inquiry—from physics to economics to trying to figure out the best place in the yard to plant tulips—is to acquire knowledge. What, then, is knowledge? What distinguishes knowledge from mere opinion? What is the best method, or methods, of acquiring knowledge? Theoretical work on the concept of knowledge is central to philosophy, because knowledge is presupposed as the target of all manner of debates, endeavors, and projects. Knowledge is valuable, both intrinsically and practically, and so it naturally occurs to us to

think about how to acquire it, and how to ensure that one's beliefs count as knowledge. To do so is to do epistemology.

A core question in epistemology is: What distinguishes knowledge from mere opinion? Can we define sufficient conditions such that, if they obtain between a subject and some object, then the subject attains knowledge of the object? Epistemologists seek to identify and characterize the most reliable paths to knowledge, and to understand what makes these paths reliable. Which beliefs count as justified, and what makes them so?

[§]

Many philosophical questions have metaphysical, semantic, and epistemological aspects—consider, for example, the questions about electrons, viruses, and persons mentioned above. In practice, it can be awfully hard to distinguish these aspects. That is, it is difficult to separate out matters of mind- and language-independent metaphysics from our beliefs about them and from the meanings of the words we use to represent them; or to separate questions about the meanings of words from metaphysical assumptions about what the words refer to and from epistemological issues about beliefs associated with the word; or to try to philosophize about knowledge in abstraction from semantic work on defining the relevant terms and metaphysical work both about what minds are and about what there is to be known. In short, it may seem like a hopeless artificial abstraction, to try to neatly separate out metaphysics from semantics from epistemology.

Nonetheless, since it is the case that many complex philosophical issues have these different metaphysical, semantic, and epistemic dimensions, exacting philosophy is marked by meticulously heeding these distinctions.¹² A key point which will loom large in the landscape to be charted herein is that metaphysical questions are questions about mind- and language-independent reality. Strictly speaking, questions of metaphysics are indifferent to human thought and talk, even if, in practice, it can be difficult to pry metaphysical issues so cleanly apart from our beliefs about them and from the meanings of the words which we use to discuss them. (For example, if we are wondering whether or not some phenomenon ought to be classified as a virus, there is a mind-independent metaphysical fact of the matter that we are trying to discern; even though various semantic and epistemological factors irreducibly shape and mold the contours of our investigation.) As will be extensively developed below, this point about mind- and language-independence underlies and grounds some important differences between necessity, on the one hand, and analyticity and *a priority*, on the other.

[§]

Appeal to a putative, controversial sort of cognition will play a key role at many junctures below. I will call it 'semantic intuition,' and it is what is

involved in grasping meanings, or understanding concepts. For example, it is via semantic intuition that one knows that all squares are four-sided, and that no grandmother is childless. (Other sorts of cognition are obviously involved in acquiring such concepts as ‘square’ and ‘grandmother’; but it is one’s source of justification for the belief that [say] all squares are four-sided that is of primary interest here, not the question of what it takes to acquire these [or indeed any] concepts. This important distinction between the cause of belief and the justification for a belief is extensively discussed in §4.1.) *How is it* that we are so sure that tomorrow scientists will not announce the discovery of a three-sided square, or of a grandmother who was not previously a mother? There surely seems to be a kind of immunity to error in this kind of case, in virtue of what these terms mean. As we will see, semantic intuition is also implicated in various kinds of philosophical issues.

One important question about semantic intuition is exactly how it relates to rational intuition, as defined above in our discussion of Plato’s problem. (Is it a kind of rational intuition? Or an alternative to rational intuition?) Further, both Quine (1951) and Kripke (1972) (among several others) have raised serious challenges for otherwise compelling presumptions about semantic intuition. In any case, semantic intuition will eventually play a role in our considered understanding of the notion of a *priority*, mapped out in Part IV.

§1.4: SUMMARY

I begin this section with a quick summary of some core distinctions and stipulations (including the introduction of a couple of terms which have not yet come up), to help to structure what follows.

[1: *a priori*] A central fissure running throughout the history of philosophy, and which is our primary focus herein, concerns whether an adequate epistemology must countenance a non-empirical (i.e., *a priori*) source of justification. Those who say YES are motivated (at least principally, among other things) by considerations of adequacy—that is, no epistemology which eschews the *a priori* can provide an adequate account of human knowledge. Those who say NO are motivated (at least principally, among other considerations) by considerations of obscurity—that is, appeal to *a priority* is too obscure to be of any (real, concrete, non-metaphorical) help in solving Plato’s (or indeed any) problem.

On some ways of understanding the rationalism/empiricism divide, the above fissure is the core point of contention in that ancient doctrinal divide; but not so on mine. On my orientation, all rationalists, and some but not all empiricists, are proponents of *a priority*. I will call empiricists who are proponents of *a priority* ‘moderate empiricists’ (and take Hume [1748], Ayer

[1936], and Boghossian [1997] to be exemplars), and empiricists who explicitly avoid any appeal to *a priori* ‘radical empiricists.’ Maddy (2000, 2007) and Devitt [2005, 2011] provide recent examples; they are typical of the species in taking Quine’s naturalism as providing motivation and inspiration.

It is a matter of some contention whether radical empiricism inevitably collapses into skepticism—for example, BonJour (1998) argues that it does, but that is not how the radical empiricists see it. (There are, of course, several distinct varieties of skepticism, distinctions between which would have to be sifted through in order to really get to the bottom of this.) One could even take the claim that to eschew the *a priori* is not (yet) to endorse skepticism to be precisely one of the core defining tenets of philosophical naturalism.

[2: *rational versus semantic intuition / acquaintance-based versus understanding-based theories of a priori*] I use the term ‘intuition’ in a way that is somewhat idiosyncratic, though by no means out of touch with its broad historical roots. ‘Intuition’ names the source of this (putative, alleged) non-empirical justification; and so (as I am using these terms) all proponents of *a priori* are ipso facto committed to some or other variety of intuition. Hence, all rationalists and some but not all empiricists hold that some such faculty is a necessary ingredient of an adequate epistemology; while radical empiricists completely dismiss intuition (at least in this sense of this diversely used term) as hopelessly obscure.

Hence, ‘intuition’ is a broad and vague notion. One important, preliminary distinction to draw within the category is that between rational intuition and semantic intuition (‘preliminary’ in that each might turn out to have multiple distinct sub-varieties). ‘Rational intuition’ is a (putative, alleged) distinctive non-empirical faculty of mind—often described as a kind of mental seeing—posited as that which is distinctively involved in cases of *a priori* knowledge. The passage from Plato’s *Phaedo* cited above provides a paradigm, seminal description of rational intuition (‘very light of the mind in her own clearness’); the classical Modern rationalists (such as Descartes [1641], Leibniz [1704]) provide others, as do some of the great logicians of the twentieth century (Russell [1912], Godel [1944, 1947]). To endorse rational intuition is to more or less adopt what I above call ‘Plato’s epistemological claim,’ in the discussion of Plato’s problem.

I will call such views ‘acquaintance-based’ theories of the *a priori*, for their positing this distinctive relation between minds and objects of *a priori* knowledge. And I will call those who posit this sort of rational intuition—or, equivalently, endorse this kind of acquaintance-based approach to *a priori*—radical rationalists. Bealer (1996, 2000), BonJour (1998), and Chudnoff (2011, 2014) are some recent defenses of varieties of radical rationalism.

Acquaintance-based theories of the *a priori* can be contrasted with ‘understanding-based’ theories, which tend to talk in terms of semantic intuition,

not in terms of rational intuition. The difference is that understanding-based theories are not committed to a distinctive (relatively mysterious) faculty of mind (or to some obscure kind of ‘mental seeing’). Rather, the idea is that all that we need for an adequate account of our universal, necessary knowledge is a thorough account of what it is to possess, grasp, or understand a concept. Semantic intuition is thus taken to be the (humble) faculty of mind involved in understanding concepts, or grasping meanings—to which we are already committed anyway, apart from consideration of Plato’s problem. Many compelling attempts to solve Plato’s problem without taking on Plato’s bold epistemic and metaphysical commitments have had semantic intuition (and the related notion of analytic truth) playing a central role.

In this region of the map, it can be hard to tell the difference between the moderate empiricists and moderate rationalists (i.e., those who identify as rationalist but tend to fashion their account of *a priori* more along the lines of semantic intuition—as opposed to acquaintance-like, rational intuition; Peacocke [2000, 2004] is a contemporary example of a self-styled moderate rationalist). Perhaps the difference is no more than that the empiricists take semantic intuition to be a categorical alternative to rational intuition, whereas the rationalists take semantic intuition to be more of a specific, relatively benign kind of rational intuition. Or perhaps, when it comes to excavating the nature and workings of semantic intuition, a more substantive divide will emerge, between moderate empiricists and moderate rationalists.

There have been times and places at which people have been tempted to distinguish these moderate empiricist/moderate rationalist camps by posing diagnostic questions about the synthetic *a priori*.¹³ However, as we will see in some depth in what follows, the analytic/synthetic distinction is in at least as much need of excavating and bolstering as the *a priori*/empirical distinction. Another, related, fairly common way to draw this divide among moderates is to ask whether *a priori* knowledge can be about (mind- and language-independent) reality, or rather whether its scope is confined to mind- and language-dependent matters (cf. Cassam [2000] for investigation). However, both the challenge of revisability and the externalist challenge cloud this way of dividing the rationalists from the moderate empiricists. These matters are extensively discussed in Part III.

[3: the constitutive *a priori*] The constitutive *a priori* is situated in this neck of the woods—either a close relative of, or a variety of, the semantic intuition, understanding-based accounts of *a priori* developed by moderate empiricists and by moderate rationalists. (Cf. the Preface above for a brief overview, and Part III below for extensive detail.) Given its explicitly Kantian flavor, it should come as no surprise that it maps onto the ancient rationalism-empiricism dividing lens in complicated ways.

The main aim of this present work is to investigate and further develop the myriad ways in which the notion of the constitutive *a priori* (i) provides an insightful, satisfactory way to approach Plato's problem, and, more generally, (ii) affords an illuminating and insightful way to understand this nexus at which semantics and the philosophy of language overlap with epistemology and metaphysics.

[§]

A central organizing theme of the present work is that there have been two main shocks to the world-order, when it comes to issues surrounding *a priori*, over the last several decades: that is, the challenge of revisability (cf. especially Quine [1951]) and the externalist challenge (cf. especially Kripke [1972]). (Aforementioned developments in the semantics of indexicality are deeply significant in themselves, and also interrelated with both of those challenges. However, while they are integrally relevant to the entire discussion (cf., e.g., §§2.2, 7.3), they do not play as central a role in the organizational structure of the book.) Here I will give a programmatic overview of the way in which my constitutive *a priori* approach incorporates these challenges; the details will be fleshed in as the work proceeds.

One epochal consequence of their collective upshot is that they reinforce a firm split between analyticity and *a priori* on the one hand, and metaphysical necessity on the other hand.¹⁴ Analytic truth and *a priori* knowledge should be understood as framework-relative and revisable—though immunity to counterexample is still preserved, in an important but limited sense. The revisability of analytic truth is entailed by the recognition that languages are organic entities which change over time. For example, on my view, it was once analytically true that whales are fish, and that there can be no such thing as a sub-atomic particle (since 'indivisible' was once constitutive of the meaning of 'atom'). However, given the subsequent changes in the relevant frameworks (and, relatedly, the changes in the meanings of the terms 'fish' and 'atom'), those things are (obviously) no longer analytically true (in any contemporary dialect with which I am familiar).

Given the close, constitutive links between linguistic meaning and mental content, many similar points will also hold of *a priori*. The root of the close links between analyticity and *a priori* lies in the fact that the (epistemic) frameworks which are constitutive of our beliefs are themselves (at least partly) constituted by (semantic) meanings. One main sub-theme herein is to defend the coherence and worth of the resultant refined notion of *a priori*, taking into account this kind conceptual evolution. I will argue that the development of the notion of the constitutive *a priori*—a pillar in common among recent work inspired by Kant, Carnap, and Wittgenstein—provides the best means to incorporate the challenge of revisability.

The externalist challenge, too, will certainly have drastic effects on our maps of this complex terrain, undermining as it does traditional conceptions of seamless, transparent connections between our concepts and their extensions. However, the range of the externalist challenge is properly limited by judicious sub-categorization within the lexicon—that is, the relevant arguments apply rather differently to ‘Aluminum is a metal’ versus ‘Widows are formerly married women whose spouse has died.’ The notion of semantic deference will turn out to be crucial for understanding the upshot of the externalist arguments: that is, to the extent that a speaker is deferential in their use of a term, the speaker does not have transparent access to the content being entertained or expressed. Further, deference is appropriate to most typical uses of natural kind terms (as well as to many sorts of uses of proper names). As Kripke (1972) shows, externalism about reference and content opens up a gap between metaphysical necessity, on the one hand, and the semantic and epistemic modalities, on the other hand.

However, there is a principled limit to the range of this externalist gap—between what constitutes the sense and what determines the reference, to borrow Frege’s (1892a) familiar terms. (Compare Loar [1991: 120]: ‘Social meanings do not deprive me of semantic autonomy when I insist on it.’) To the extent that the speaker exhibits semantic autonomy, and is not deferential, then the speaker has transparent access to the content being entertained or expressed, and the externalist gap does not open up. Furthermore, there are vast ranges of the lexicon (e.g., ‘and,’ ‘hunter,’ ‘bachelor,’ ‘triangle,’ ‘chair,’ ‘ball,’ ‘spoon,’ etc.) concerning which autonomy is more appropriate than deference, and hence to which the externalist arguments have little applicability.¹⁵

Many traditional tenets will be preserved in the resulting overall maps of the terrain sketched in Part IV. There are (framework-relative) analytic truths, and they are one and all knowable *a priori*. The converse does not hold though—one of the many important morals of our refined understanding of indexicality is that it provides relatively uncontroversial instances of things which are justified *a priori* but not analytically true (e.g., ‘I am conscious’). Metaphysical necessity will turn out to be a completely different matter than either *a priority* or analyticity, particularly when it comes to the aforementioned case of deferentially used natural kind terms. (For example, ‘Water is H₂O’ might be necessarily true without being either analytic or *a priori*.) There is nothing framework-relative or revisable about necessary truths; and hence many traditional habits of transitioning between necessity and either of the other two concepts—in either direction—have been undermined.

[§]

Onwards then with developing the case that the constitutive *a priori* can help us to reach a richer understanding of the relations between these

metaphysical, epistemic and semantic varieties of immunity to counterexample, and ultimately of Plato's problem.

NOTES

1. This notion of '*a priori*' has been the predominant one at least since Kant (1781). In some earlier usage, the '*a priori*'/'*a posteriori*' contrast depended on whether reasoning went from causes to effects or from effects to causes; and so Descartes' (1641) '*a priori*', for example, is not exactly Kant's or ours—though there is significant overlap. Clearly, though whatever it might have been called at the time, the notion of non-empirical justification was clearly involved in the work of Plato, Descartes, and most other major pre-Kantian philosophers.

According to Burge (2000: 13), Leibniz was the first to define the now standard notion of *a priority*, which is commonly associated with Kant. Pap's (1958) thorough analysis of this notion of *a priority* also begins with Leibniz. For a good overview of historical issues pertaining to *a priority*, cf. Hanson & Hunter (1993); and for a good recent statement of the centrality of the *a priori* to a variety of philosophical debates, cf. Peacocke (2006).

2. It has been objected that [1] is a bad example because it is subject to counterexample in cases of co-ownership. I do not think that this putative counterexample succeeds—if you and I co-own something, to whatever degree, I could only possibly steal your part of it, not the part that is already mine. However, this objection does help to show just how difficult it is to come up with an uncontroversial example of *a priority*.

The notion of the constitutive *a priori* developed in Parts III and IV will help to explain why this is so—that is, while *a priority* is coherent and significant, it is framework-relative, all frameworks are more or less constantly under revision, and many distinct frameworks can be simultaneously pertinent to any given question or issue.

3. There is a growing body of literature which distinguishes between such claims as that *a priori* knowledge, or analytic truth, is (i) universal, (ii) eternal, (iii) indefeasible, etc. (Cf. Casullo [2003] for discussion of the case of *a priority*, and Russell [2008] for discussion of various distinct senses of analyticity.) We will get into such fine distinctions in Part II.

4. See Kripke (1972: 38–9) for a nice brief statement of the grounds for thinking that all and only necessary truths can be known *a priori*; this point is also further developed below in §1.2. As for 'analytic', there are also considerable grounds for tight connections between it and each of these other two concepts. Indeed, historically and conceptually, much of the interest in analyticity has stemmed from its promise to illuminate the concepts of necessity and *a priority*.

5. For discussion of Hobbes' views see Munsat (1971: 19–20); Locke's discussion occurs at (1690: Bk. IV, sect. 8).

6. Kant's (1781) famous attempt to distinguish '*a priori*' from 'analytic' is one notable effort to the contrary. However, as we will see, Kant's efforts at establishing an analytic/synthetic distinction have not withstood scrutiny. Canonical critical discussions include Frege (1884), Quine (1951).

7. In addition to Quine, other seminal work pertaining to what I am calling the challenge of revisability (cf. Part III for extensive discussion) was done by Mates (1950), White (1950), and Goodman (1952).

8. Kripke (1972) is the most thorough and influential source, when it comes to the externalist challenge (cf. Part III for extensive discussion). Other important contributions include Donnellan (1970), Putnam (1975), Burge (1979), and Kaplan (1989).

9. To be sure, this does presuppose a certain degree of realism, which some might find controversial. The relevant degree of realism is nicely articulated by Stalnaker (1984: ix) as follows: ‘The world is the way it is independently of our conceptions of it, and the object of inquiry is to find out the way the world is.’ I will not offer any arguments, to try to convince those who cannot tolerate even that degree of realism; though it is important to note how much more cautious and conservative it is than are many of the notions of ‘metaphysical realism’ floated by philosophers who seek to criticize or reject realism (e.g., Putnam [1992]).

10. Note that we are not talking about the symbols, but about what they signify: We might have spoken a language in which the sentence ‘ $2 + 2 = 4$ ’ means that seven is even, or that Bob’s your uncle; but that would not be a situation in which two of anything added to two of anything would fail to be equal to four of them. There is further discussion of this important point in §2.1.

11. Experimental philosophy is a recent development within this naturalistic tradition; and there is a growing body of literature dedicated to sorting out whether it has any deep or drastic relevance to the *a priori*. I will not pursue that question herein. Cf. Weinberg (2014) for an insightful account of how the tools and methods of experimental philosophy can be seen to supplement, rather than undermine or overthrow, traditional *a priori* tools and methods. As should become clear by Part III (cf., e.g., the discussion of Wittgenstein [1969] in §6.1), this idea that the *a priori* and the empirical are not discrete containers, but are rather commonly blended together, in complex ways, in the course of inquiry, is also a guiding tenet of the constitutive *a priori* orientation.

12. For example, one cannot sufficiently justify any metaphysical conclusion using only epistemic or semantic premises. To illustrate, consider what is known as a ‘fallacious appeal to ignorance’—for example, no one can prove that P is false; therefore P. Among the problems with this pattern of reasoning is that its premise is purely an epistemic one (about who knows what), while its conclusion is a metaphysical claim to which such epistemic considerations, while relevant, cannot be considered decisive. Many other such issues pertaining to transgressing borders between these different sub-disciplines will be discussed below. (Cf., e.g., note 14.)

13. Indeed, many proponents (e.g., Bonjour [1998]) and opponents (e.g., Aune [2008]) of rationalism alike hold that rationalism is an interesting and significant position only if Kant is right that there are synthetic *a priori* judgments. Otherwise why fuss about rational intuition, if it is unable to amplify and add to the store of human knowledge?

14. Hence, for example, on the view defended in Parts III and IV, conceivability absolutely and unequivocally does not entail metaphysical possibility. Conceivability is a main source of evidence about possibility, but it is deeply defeasible evidence.

Conceivability tells us—first and foremost—about concepts, not about their referents or extensions; and it is the extensions of concepts which are relevant to questions of metaphysical modality.

15. This lack of an externalist gap between what is constitutive of the sense and what determines the reference is, significantly, also evident in the case of the so-called ‘*a priori* sciences’ of logic and mathematics. More on this in Part III (especially §§6.2 and 6.4–5).

Chapter 2

Further Preliminaries

This chapter is dedicated to increasing some prerequisite technical precision when it comes to certain points about concepts, thought, and language which will be important for our ongoing project. In particular, I mentioned in §1.3 that there have been many leaps forward in semantics and the philosophy of language in twentieth-century philosophy, which are pertinent to a refined understanding of *a priori*. It will prove worthwhile to chart some of this territory, before turning to critical work on the varieties of immunity to counterexample, in Part II.

§2.1: SOME IMPORTANT TERMS AND CONCEPTS

First up: There is considerable diversity in the literature, when it comes to the question: To what, precisely, should the terms ‘necessary,’ ‘analytic,’ and ‘*a priori*’ be attributed? (This is especially true of the term ‘*a priori*,’ as we will see in §4.1. Truths, propositions, concepts, inferences, knowledge, and justification are among the many things which are commonly labeled *a priori*.) I will begin this section with discussion of two common and central candidates: truths and propositions.

‘Necessary truth,’ ‘analytic truth’ and ‘*a priori* truth’ are all common locutions in philosophy. All three of these terms target special categories of truth, truths not limited to specific people, times, or places. Like many much-discussed philosophical concepts, ‘truth’ tends to be employed rather variously. One hears that truth is culturally constructed, that truth is power, truth is historical, truth is whatever the rich say it is, and so on. One is apt to encounter the sentiment that you have your notion of truth, I have mine, and no one is in a position to comparatively evaluate them.

Although not without motivation, these sentiments are, as stated, rather sloppy and misleading. ‘Truth’ in the above sentiments stands not for the content of the concept *truth*, but rather for the set of things that are taken to be true. It is the set of things that are taken to be true that varies from culture to culture, that the rich and powerful have a disproportionate say in dictating, and so on. You may well be free to take as true whatever you see fit, and no one else has authority to decree otherwise, but none of that touches the one constant content of the concept of truth. It would be rash to conclude from these wide variations in what is taken to be true that the content of the concept of truth varies widely. Quite the contrary: disagreements about which things are true depend upon agreement about what the term ‘truth’ means.

In this book I will adopt a deflationary stance concerning the content of the concept ‘truth.’¹ Strictly speaking, truth is a simple, elementary concept. One cannot learn a language, or tell a lie, without manifesting a grasp on it. ‘Truth’ names a basic relation between linguistic expressions and what they are about, on which three-year-olds have a decent purchase, and with which we are all thoroughly competent long before adolescence. There are all manner of difficult philosophical and political questions about who gets to say what is to be counted as true, about how what is taken to be true varies across perspectives, and so on; but these debates presuppose, rather than call into question, the one constant content of the concept of truth. In the *Metaphysics*, Aristotle makes the earth-shaking assertion that truth consists of saying of what is that it is, and of saying of what is not that it is not. Two and a half millennia later, after the rise and fall of some much more sophisticated theories of truth, that’s about where matters stand. As Quine (1992: 82) puts it: ‘One who puzzles over the adjective “true” should puzzle rather over the sentences to which he ascribes it. “True” is transparent.’

One conclusion that is fairly well supported by the many failed attempts to give a non-circular, substantive definition of truth is that it probably cannot be done. Truth cannot be reduced to any other terms. This elementary status should not be all that surprising, and is hardly a defect—very few concepts can be reduced to others without remainder. (Potential examples of concepts that can be so reduced include ‘snowball’ [= made of snow + shaped like a ball] or perhaps ‘bachelor’ [= unmarried + man]). It is now widely recognized that, in this respect, these concepts are in the minority. Just try to break down ‘freedom,’ ‘happiness,’ ‘good,’ or ‘art,’ say, into self-evident, discrete factors! In general, this inability to factor out discrete and exhaustive constituents of a concept is not a sufficient reason to hold that that renders the concept suspect, or second-rate.²

[§]

I depart from Quine (1992: 82) in taking the things for which the question of truth or falsity arises to be not sentences, but rather the propositions that

sentences are used to express (in context). So, in calling something a ‘truth,’ I mean that the proposition expressed by the given sentence is true. The term ‘proposition’ is introduced to have one to denote what is in common among, for example, uses (in context) of the following sentences:

1. Aristotle is now sitting.
2. I am now sitting. (said by Aristotle)
3. You were sitting, yesterday at this time. (said to Aristotle, in 24 hours)
4. *Aristote est assis, maintenant.* (said in Quebec)

[1]–[4] describe the same situation, express the same information, or say exactly the same thing. Henceforth, I’ll say that they express the same proposition. Propositions are more abstract than sentences—the same proposition can not only be believed or doubted by different speakers, but can be expressed by means of different sentences, in different languages, and so on.

So, first and foremost, the term ‘proposition’ is a useful one in this kind of inquiry, because it allows one to categorize various thoughts and utterances into equivalence classes, based on their representational properties, or information content. Over and above that point about usefulness, though, some such abstract term as ‘proposition’ is essential in this kind of inquiry, because it is only at this level of information content that questions of truth or falsity—and, a fortiori, question about immunity to counterexample—arise.

Propositions are means of categorizing attitudes and assertions into equivalence classes. Within the philosophy of language, there are many complex debates about the precise nature, contents, and individuation conditions of propositions. It would be tangential to our present project to delve very deeply into these waters, but some pertinent refinements will be explored in this chapter. For example, prior to—and deeply pertinent to—the question of whether ‘Water is H₂O’ is necessary, analytic, or *a priori* is the question of how precisely to individuate the content of the proposition expressed; and on that question, there are many distinct options.³

In addition to ‘sentence’ and ‘proposition,’ I also sometimes avail of the notion of ‘statement’ below—where a statement is a particular dated use of a sentence to express a proposition.

[§]

I next consider some miscellaneous notes on some terms in the neighborhood of our core notions of necessity, analyticity, and *a priority* which have some currency, but which will not play a role in this book.

One such term is ‘epistemically necessary,’ which is used to mean something like ‘necessitated by one’s antecedent beliefs’ or ‘entailed by things one holds true.’ As I use the terms in this book, ‘epistemic necessity’ is

not well-formed—since ‘necessity’ is not an epistemological concept. To speak of such explicitly epistemic concerns, I will use explicitly epistemic vocabulary.

Consider next the common terms ‘conceivable’ and ‘conceptual truth.’ As above, given my above setup, this is a kind of hybrid that vaguely straddles distinct philosophical sub-terrains of necessity, analyticity, or *a priori*. Most specifically: I take ‘conceptual truth’ to typically mean a certain semantic/epistemic hybrid—that is, ‘knowable *a priori* because it is analytically true.’ In any case, by Part IV, I will have come down firmly against the metaphysical import of anything like ‘conceptual truth’; for conceivability tells us about our concepts, as opposed to about the mind-independent objects which make up the extension of (most of) our concepts.

Another such term is ‘logically necessary,’ which may or may not be broader than metaphysical necessity (just as metaphysical necessity is broader than physical⁴). Logical possibility is supposed to include everything conceivable, anything that is not precluded by the laws of logic, where these are understood to be distinct from, and perhaps more inclusive than, the laws of metaphysics. I shall not make use of this notion here, either. I am not sure that this distinction between metaphysical and logical laws could be drawn in a clear and comprehensive way, let alone what pertinence this notion of logical necessity would have for our ongoing discussion anyway, in addition to the three notions of metaphysical necessity, analytic truth, and *a priori* knowledge. Anyway, nothing which follows either presupposes or further investigates any distinction between logical and metaphysical necessity.

I will have little to say that is directly about certain other important, nearby notions of logical truth, entailment, validity, etc. First and foremost, these notions pertain to relations between propositions; whereas my primary interest here lies more at the level of properties of propositions. While logic and the philosophy of logic are relevant to parts of our ongoing discussion (cf., e.g., §6.2), for the most part, they will not be directly engaged in a sustained way. (And, even there, it is a case study in epistemology, as opposed to logic or the philosophy of logic per se, which is our focus.)

§2.2: THE ARBITRARY NATURE OF LINGUISTIC CONVENTIONS

It is important to examine at a bit more length the arbitrary nature of linguistic conventions. Let us start with some candidates for necessity, analyticity, or *a priori*:

1. $2 + 2 = 4$.
2. One cannot steal one's own property.
3. All bachelors are unmarried men.

For present purposes, it is crucial to hold linguistic conventions fixed, because what is at issue is not the identity or nature of the symbols, but rather what the symbols express. The claim that [1] is necessary, or that [2] is knowable *a priori*, say, is not in tension with the fact that our linguistic conventions are, in some sense, arbitrary. We came up against this matter in §1.2, in considering the claim that [1] is necessary: there it was stated that we might have spoken a language in which the sentence '2+2=4' meant something different, but that is not relevant to the modal status of [1]. The claim that [1] is necessary would not be falsified if that sequence of squiggles were to be put to some other use; rather, it would be falsified only if—regardless of which kinds of objects you consider, or in which order you count them—two of anything added to two of anything failed to yield four of them. This distinction between symbols and what they are used to express is absolutely fundamental here, and so I should elaborate it a bit further. (As we will soon see below, and again in §3.3, some skeptical arguments about the very idea of immunity to counterexample founder on cautious distinctions in this neighborhood.)

Given the distinction between sentences and propositions, the point is that there is nothing necessary about *sentences*; it is only at the level of *propositions* that the question of necessity arises. (This might seem so basic as to not need pointing out; but especially when we come to certain arguments against the coherence and worth of analyticity, these basic points will prove their worth.) The claim that [1] is necessary, or that [2] is knowable *a priori*, say, is consistent with the fact that our linguistic conventions are, in some sense, arbitrary. The claim is not that these sentences could not possibly have been used to mean something false, or that they might not express contingencies in another possible language. It is rather that, holding linguistic conventions fixed, what [1] means could not possibly be false, and that [2] can be known to be true without performing experiments or taking polls, but rather just by reasoning things through. In considering these claims of immunity to counterexample, we must first fix on the proposition expressed by the sentence; and then we consider whether *that* could possibly be false (in the case of questions of necessity), whether it is true in virtue of meaning (for questions of analyticity), or whether it can be known to be true without empirical investigation (in the case of *a priori*).

The relation between any linguistic expression and its meaning is somewhat arbitrary—that is, in no case is it necessary that any particular sound or symbol (e.g., '2,' 'no,' or whatever) has the particular meaning that it does.

These symbol-sound-meaning conventions are historical accidents, contingent facts set by the evolution of our language. As a consequence, there is something arbitrary about the relation between sentences and the propositions that they express—since historical accidents play a role in the meanings of all the constituent parts of, say, ‘One cannot steal one’s own property,’ the link between sentence and proposition is also a complex conventional historical accident. So, as linguistic conventions vary, which proposition is expressed by a given sentence will vary. If ‘steal’ meant what ‘sell’ actually does, then [2] would express a different proposition; if ‘property’ meant what ‘orangutan’ actually does, then [2] would express a different proposition; and so on.

Hence, if we did not hold linguistic conventions fixed, all would be (trivially) contingent—it is in no case necessary that any given sentence express the particular proposition that it does. Crucially, though, necessity, analyticity, and *a priori* have nothing to do with alternative possible linguistic conventions. Rather, they have to do with the meaning that is expressed with a given utterance or inscription of a sentence. As Kripke (1972: 77) puts it:

One doesn’t say ‘two plus two equals four’ is contingent because people might have spoken a language in which ‘two plus two equals four’ meant that seven is even.

To bring up alternative possible meanings, in adjudicating our questions (of *a priori*, necessity, and analyticity), is to change the subject.

The issue of precisely how we determine which meaning is expressed with a given sentence is quite complex. The crucial point for now is that claims of analyticity, like necessity and *a priori*, are not claims about other possible linguistic conventions, and so cannot be undermined by appeal to them.⁵

[§]

These above distinctions pertaining to linguistic conventions are most crucial when it comes to certain lines of attack on the notion of analytic truth. Consider for example the following possible interpretation of ‘analytic truth’—that is, way to understand exactly what ‘truth in virtue of meaning’ comes to—described by Boghossian (1997: 336): ‘Our meaning *p* by *S* makes it the case that *p*,’ where ‘*S*’ stands for a sentence, and ‘*p*’ stands for the proposition it expresses.⁶ This sort of view would take our meaning-conventions to have some supernatural magical potency. It reads the ‘in virtue of’ in ‘truth in virtue of meaning’ as a *causal* connection. The idea seems to be that our meaning-conventions have the power to move things about in mind- and language-independent reality, to change the nature of the things about which we think and talk.

Boghossian (1997) clearly rejects this conception of analyticity. I, too, want nothing to do with this strange view. (As far as I can tell, no one ever

held this view.) The claim that [1] or [2], say, is analytic is not the claim that linguistic conventions *caused* the relevant mind- and language-independent facts to fall into place.

[1] $2 + 2 = 4$.

[2] Murder is wrong.

Rather, if our meaning-conventions determine that ‘+’ and ‘four’ (or that ‘murder’ and ‘wrong’) express certain concepts, and if it is plausible that the relations between these concepts are immune to counterexample, then there will be statements such that grasp of their meaning is sufficient for recognition of their truth.

Here we see the importance of the above distinctions between conventional links between S and p and truth-conditions of p. First, it is a contingent conventional accident that such symbols as ‘square’ or ‘four-sided’ are used to express the concepts that they in fact express; second, it still yet may be immune to counterexample that all squares are four-sided. Given that some relations among concepts are not vulnerable to refutation by contingent happenstance, it would be awfully odd if we were somehow barred from making statements whose role is to express such exceptionless relations.

There is, however, an incontrovertible point in the neighborhood of the above straw conception of analyticity. The straw conception has it that:

[S] Our meaning p by S makes it the case that p.

In contrast, the distinct, more refined point is that:

[Si] Our meaning p by S makes it the case that, if p is impossible, then what S expresses is false in every possible circumstance.

[Sii] Our meaning p by S makes it the case that, if p is contingent, then what S expresses is true in some circumstances and false in some circumstances.

[Siii] Our meaning p by S makes it the case that, if p is necessary, then what S expresses is true in every possible circumstance.

The key difference is that [Si–iii] explicitly separate out two very different factors which are relevant to the question of a statement’s truth-value—that is, first, there is the conventional link between S and p; and second, there is p’s truth-condition.

Again, if we did not hold linguistic conventions fixed, in the course of our present inquiry, all would be trivially contingent—since all links between sentences and propositions are contingent. Crucially, though, necessity,

analyticity, and *a priori* have nothing to do with alternative possible linguistic conventions. Those notions concern what statements express, not what they might have expressed. To bring up questions about alternative meanings, in these kinds of discussion, is to change the subject.

[§]

Perhaps not surprisingly, given his naturalistic agenda, Quine was prone to run roughshod over this distinction between the conventional link between S and p, on the one hand, and the truth-conditions of p, on the other. Consider, for example:

The statement ‘Brutus killed Caesar’ would be false if the world had been different in certain ways, but it would also be false if ‘killed’ had the sense of ‘begat.’ (1951: 36)

There are enormous relevant differences between the two conditions cited by Quine. If the world had been different in certain ways, then Brutus might not have killed Caesar, but if ‘killed’ had the sense of ‘begat,’ then the sentence ‘Brutus killed Caesar’ would express a rather different proposition. If we uncritically sweep the contingencies of linguistic conventions together with the contingencies of language-independent facts, in this way, then the trivial, misleading pseudo-conclusion that all is contingent quickly follows.

To the contrary, whereas Quine treats these two separate truth-value-determining factors (i.e., the conventional link between S and p, and the truth-conditions of p) as if they were of the same sort, these factors are operative at different levels. Linguistic conventions have to be settled first, and held fast, so that we can fix on a particular p; after which point linguistic conventions are entirely irrelevant to modal questions about p. On this point, Stalnaker articulates a clear advance beyond Quine (1951):

[W]hen a statement is made, two things go into determining whether it is true or false. First, what did the statement say: what proposition was asserted? Second, what is the world like; does what was said correspond to it? (1972: 177)

Stalnaker’s first question is the conventional one, whereas our ongoing focal questions (about necessity, analyticity, and *a priori*) only pertain at the level of the second question, subsequent to settling on answers to the first question. (We’ll pick up and build on this shortly, in the next section.)

So, insofar as Quine’s point was merely that ‘in general, the truth of statements depends obviously on both language and upon extralinguistic fact’ (1951: 41), then well and good—I completely agree. However, that by no means entails that there is no such thing as immunity to counterexample in virtue of meaning:

- [1] No grandmother is childless.
 [2] Every even number is divisible by two.

While it is awfully hard to draw a firm boundary around the set of analytic truths, there clearly are paradigm cases. What would you say to someone who thought it an open question whether all grandmothers are women who are or were previously mothers, thought that we should do a poll or conduct experiments before we could confidently claim that [2] is true? The conclusion that they simply do not understand the meanings of the relevant terms is compelling if not inevitable; and it will be grounded within a sophisticated framework, below in Parts III and IV.

There is nothing obscure or magical at work here; the only requisite ingredients are (i) conventional relations between expressions and properties, and (ii) necessary relations among properties. (We will get into putative analytic truths which are not necessary in due course.) Of course, both metaphysical necessity and the determinacy of meaning raise hard questions. The present point is just that anyone who accepts both the notions of necessity and of meaning has already purchased all the ingredients for analytic truth; since analytic truth is widely held to be such a far cry more contentious than either metaphysical necessity or the determinacy of meaning, the point is worth making here.

§2.3: MEANING, EXTENSION, INDEXICALITY, MODALITY

Up to this point I have occasionally availed of a fairly uncritical, intuitive notion of ‘meaning,’ and have said just a little bit about the indispensable notion of a ‘proposition’ (i.e., basically, equivalence classes of sentence-sized meanings). The next business here is to increase the technical precision when it comes to these two notions, particularly when it comes to the ways in which developments in the semantics of indexicality and of modality (among other related developments) have altered the terrain.

The first basic refinement to ‘meaning’ is as follows: I use the term ‘extension’ to designate the thing or set of things to which a given expression refers, or correctly applies. For example, the extension of ‘persimmon’ is a certain bunch of fruit; the extension of ‘the richest woman in Europe’ is the person who satisfies the condition specified; and the extension of ‘midnight blue’ might be taken to be the scattered and diverse collection of spatial regions which instance the relevant shade.

Although we sometimes say such things as “‘Persimmon’ means *that kind of fruit right there,*” or “‘Midnight blue’ means *that shade right there,*” I

will never use ‘meaning’ in this way here. The term ‘extension’ will be used exclusively for this sort of job.

‘Meaning,’ in contrast, will be used to designate the significance or connotation of an expression. The meaning is what a dictionary aims to record; it is the conceptual condition that must be grasped in order to understand the expression. So the extension of ‘persimmon’ is a bunch of fruit, but its meaning is the conceptual condition that picks out all and only that fruit. To know the meaning is to grasp that condition. This meaning/extension distinction is starkly illustrated by the case of definite descriptions, such as ‘the richest woman in Europe.’ If one understands all the constituent words, and is competent with the grammar of English, then one grasps its meaning; but knowing who its extension is goes beyond linguistic competence. One must actually know which person happens to satisfy the relevant condition. It is easy to generate definite descriptions whose meaning we all grasp but whose extension probably none of us do (e.g., the biggest fish in the Indian Ocean, the tallest left-handed Mexican in Finland).

There is obviously a close connection between meaning and extension—that is, the meaning specifies the conditions that must be satisfied in order to qualify for membership in the extension. Meaning, in context, imposes constraints on what can count as the extension. To know the meaning is, very roughly, to be able to identify the extension (in normal contexts), to distinguish between the extension and its complement.

Now, a part of why we need to be as clear as possible about such semantic matters is that the notion of synonymy (i.e., sameness of meaning) is crucially important to, first and foremost, analyticity, and, subsequently, to *a priori* as well. Sameness of extension is clearly a necessary condition for synonymy.⁷ If, say, someone thought that ‘persimmon’ and ‘pomegranate’ were synonyms, but then became convinced that there was an object to which only one of those terms correctly applied, then that person would have to reject the previously held hypothesis about the terms’ synonymy.

Throughout the history of theorizing about logic and language, many have also wanted to hold that sameness of extension is a sufficient condition for synonymy. In that case, meanings would be relatively easy to incorporate into a naturalistic world-order. Meanings would obey set-theoretic axioms of extensionality, as well as other intuitive principles of extensional logic. (Key ones include the substitutivity of co-referential terms *salva veritate* [i.e., if $a=b$ and Fa , then Fb] and existential generalization [i.e., if Fa , then $\exists xFx$].) Talk of meaning could be translated into, or transparently reduced to, well-understood set theory. The meaning could be identified with the extension; and hence otherwise obscure questions about sameness of meaning would be transformed into relatively clear questions about whether distinct sets do or do not have the same members.

However, meanings do not seem to be so simple. It appears that distinct expressions may differ in meaning even while being co-extensive. Classic cases include 'renate' (or: organism with kidneys) and 'cordate' (or: organism with a heart). While clearly the conditions for membership in these sets differ, these conditions determine exactly the same members. That is, as a matter of fact, all and only renates are cordates. Hence, at least *prima facie*, these terms are co-extensive but not synonymous. Other related types of examples include distinct terms with null extension, which do not seem to be thereby synonymous (e.g., 'phlogiston' vs. 'perpetual motion machine' vs. 'frictionless plane,' and so on). Frege famously argued that the same phenomenon occurs in the case of proper names (e.g., 'Hesperus' vs. 'Phosphorus'), but while this claim that such co-extensive names differ in meaning is plausible and has been influential, it is too controversial to rest much upon, without the kind of extensive further analysis of the semantics of proper names which will not be undertaken in this present work.

In any case, I take the renate/cordate cases to at least strongly suggest that more than just extension seems to be involved, when it comes to questions of meaning. Hence, talk of meaning cannot obviously be translated or reduced to talk of extensions. This situation is deplored by naturalists such as Quine (1956), who laments that meanings are 'obscure ... creatures of darkness'. A comprehensive theoretical account of human beings must appeal to and incorporate a theory of meaning, but as of yet meanings have not been naturalistically, seamlessly incorporated into a scientific world-order.⁸

[§]

Next to move from individual-expression-meanings to sentence-sized meanings (i.e., propositions). In what Kaplan (1975) calls 'the Golden Age of Pure Semantics,' the links between individual atomic meanings and propositional molecules were taken to be simple and transparent. Every independently significant linguistic expression was taken to have a context-independent meaning. Semantic competence requires pairing expressions with their meanings. When expressions are put together into an intelligible sentence, those meanings and their mode of composition determine the content of the proposition expressed. Propositions themselves were then complex sentence-sized meanings which constitute (among other things) the truth-conditions of an assertoric use of the sentence, and that which must be grasped in order to count as understanding the meaning of the sentence. Some would even go so far as to characterize the semantic enterprise itself via this end of pairing up well-formed sentences with propositions expressed (cf. Lewis [1975]).

The Golden Age is long gone now. (Kaplan himself cites Carnap [1947] as its zenith.) Various different factors have served to show that the relations

between atomic meanings and propositional molecules is vastly more complex than that. Key developments here include developments in the semantics of indexicality and modality.⁹ These developments are very important, for our purposes, since: [1] they complicate various senses in which something might be taken to be immune to counterexample, and [2] they necessitate distinctions which are in any case crucial for an adequate degree of semantic precision.

Indexicals are linguistic expressions whose extension shifts from use to use (e.g., ‘she,’ ‘today,’ ‘here,’ ‘this’). Here we have a common kind of context-sensitivity in which sameness of meaning is compatible with differences in extension. For example, various utterances of ‘Today is rainy’ or ‘She is German’ (in different contexts) can clearly express truth-conditionally distinct propositions. The semantics of indexicality studies the factors which determine the contents of propositions expressed in such cases. Indexicality is clearly in tension with the Golden Age conception of the relations between sentences and propositions; one important question is whether these indexical cases are circumscribed oddballs, or else the thin end of the wedge.¹⁰

Modality poses a similar but very different challenge to the Golden Age conception. Modality is an ancient area of philosophical inquiry that centers on concepts like *possibility*, *contingency*, and *necessity* (or, in their more common guises, *might*, *can*, and *must*). Drawing modal distinctions can be seen as a matter of conceiving of and reasoning about non-actual but possible contexts of evaluation. Thus, insofar as we would concede that *although Neil Young is not in fact the prime minister of Canada in 2017, he might have been*, that is to say that there is a consistent, coherent non-actual situation in which different historical accidents befell Neil Young and he became a successful Canadian politician instead of a folk-rock icon. In contrast, insofar as we would also concede that *Neil Young could not have been a coffee cup*, that is to say that there is no such consistent coherent scenario in which historical accidents could have conspired to make our actual Neil Young into such an object. Hence, the study of modality involves consideration of distinct non-actual contexts of evaluation.

Now, recall that on the Golden Age conception of things, interchange of co-extensive parts ought to preserve truth-conditions. If all renates are cordates, then ‘All John’s pets are renates’ is true iff ‘All John’s pets are cordates’ is. However, modality provides one fairly clear case in which this Golden Age tenet is subject to counterexample. Suppose actually true anything of the form ‘a = the F’ (e.g., Justin Trudeau is the prime minister of Canada, or Adele is the richest woman in Europe); even still, ‘a is G’ and ‘The F is G’ nonetheless differ in their truth-conditions across various possible contexts of evaluation (e.g., ‘Justin Trudeau is charming’ vs. ‘The Prime Minister of Canada is charming’). There clearly are consistent coherent non-actual situations in

which Justin Trudeau is charming but the prime minister of Canada is not, for example, or in which Adele is rolling in the deep but the richest woman in Europe is not.

Hence, the kind of consideration of non-actual contexts of evaluation involved in the study of modality points to another important sort of refinement to the Golden Age. Co-extensiveness of parts is no guarantee of sameness of truth-condition across contexts of evaluation. This is another major complication to the relations between linguistic meaning and propositions expressed, which will be pertinent in what follows. (It is also another reason to be wary of the reduction of meaning to extension.)

It is interesting to consider these complications necessitated by indexicality and modality against the backdrop of a quote from Stalnaker cited in the last section:

[W]hen a statement is made, two things go into determining whether it is true or false. First, what did the statement say: what proposition was asserted? Second, what is the world like; does what was said correspond to it? (1972: 177)

Indexicality is complexity about the relation between linguistic meaning and the first level—factors in the context of utterance can affect which proposition gets expressed. In contrast, modality pertains to the relation between proposition expressed and its conditions of satisfaction, and hence pertains to the second level. Factors in the context of evaluation can affect whether the proposition that is expressed is true or false.

Generally, investigation into indexicality and modality lead to the development of multi-dimensional semantic frameworks in which the effects of the context of utterance on content of the proposition semantically expressed can be neatly distinguished from the effects of the context of evaluation on whether that proposition is true or false (cf. Kaplan [1989], Stalnaker [2001], Chalmers [2006]). These refinements to the relations between the meanings of linguistic expressions and the propositions which are expressed when they are used in context will be crucial to subsequent discussions of immunity to counterexample—first in §3.3, and then throughout Parts III and IV. There will be many significant further appearances by both indexicality and modality.

[§]

Now to move onto propositions per se, that is, the complex meanings of complete sentences. Due in large part to increasing sophistication in our understanding of such phenomena as indexicality and modality, we now distinguish several varieties of sentence-sized semantic entity, each of which is suited to some but not all of the jobs traditionally associated with the term

‘proposition.’ In the course of progress, Golden-Age, canonical propositions (as Frege or Russell conceived of them) have been methodically pulled apart.

The phenomenon of indexicality provides a relatively simple way to illustrate this point. Given a use of ‘I am here now’ (in context), there are (at least) two different sentence-sized semantic entities which must be distinguished, both of which are plausibly thought of as (in some sense or other) semantically expressed by the utterance. There is what Kaplan (1989) calls the ‘character’ (i.e., the information that is identifiable independently of the context of utterance), and then there is what he calls the ‘content’ (i.e., the information that issues once the character becomes saturated by the relevant features of the context of utterance).

To the extent that indexicality is thoroughly prevalent, important differences emerge between character and content. For present purposes, consider just two of the proposition’s main jobs—to constitute that the grasp of which constitutes linguistic competence, and to specify the truth-conditions. As for ‘linguistic competence,’ insofar as we demand of propositions that they be that which is grasped by one who understands the sentence (whether or not they can ascertain all pertinent features of the context of utterance) then it is sentential characters which fit the bill. (Common examples for making this point include receiving a postcard in the mail, with the author’s name and location illegible, that says ‘I am having a great time here’. Linguistic competence gives you the character, but you need to solve for WHO and WHERE to get to the content.) In contrast, insofar as propositions are the bearers of truth-conditions, then here contents are much better suited.¹¹

In this respect, indexicality is just the thin end of the wedge. Things get more and more complicated for poor old golden-age propositions once we recognize the prevalence of context-sensitivity, semantic underdetermination, etc. Even further, much has transpired since the 1970s, when it comes to the various semantic entities associated with sentences; such that by this point there are several other alleged denizens in what Taylor (2007) calls ‘the sub-syntactic basement of the language’—unarticulated constituents, hidden variables, non-classical relativistic parameters, etc. It is not altogether clear exactly how such posits should be taken to affect the individuation conditions for propositions¹²; though again here we are straying into technical issues within the philosophy of language which lie beyond the scope of our present inquiry.

[§]

For the case of sub-sentential expressions, I am going to stick with the terms ‘meaning’/‘extension,’ as opposed to taking on Kaplan’s ‘character’/‘content.’ There is broad overlap between these contrastive pairs, over a range of paradigm cases; but ultimately Kaplan’s terms are integrally connected to his

technical artifice in a way that may not well-suit the scope of this project. (When you pick terms as common and varied as ‘meaning’/‘extension,’ people pay attention to your exact stipulations as to how they are to be used; but when you talk in terms of ‘character’/‘content,’ the boundary between novel research and Kaplanian exegesis gets blurrier.)

One possible drawback to ‘character’/‘content’ is that while it clearly applies to indexical pronouns, it is less clear or obvious how to apply it to various other sorts of expression. When it comes to ‘the richest woman in Europe,’ for example, what exactly is its content? A person? An identifying condition? Something else? The general run of quantified noun phrases (‘a man,’ ‘three thugs,’ ‘everyone’) is illustrative of the possible cases for which (as opposed to ‘she’ or ‘today,’ used in context) it is not clear exactly what the content should be taken to be. From the other direction (and again as opposed to ‘she’ or ‘today,’ used in context), it is not clear what the character of proper names or natural kinds terms, for example (e.g., ‘Aristotle,’ ‘aluminum’), should be taken to be. So, in contrast to the characters versus contents of indexical expressions, as stipulated above, the ‘meaning’/‘extension’ pair has a broader range.

A second reason not to go with ‘character’/‘content’ is that Kaplan’s talk of characters as functions from contexts to contents, though helpful in many respects, has unfortunately helped to engender some oversimplified pictures of the relations between uses of linguistic expressions and the propositions thereby expressed.¹³ Only in the case of ‘I’ (sometimes called an ‘automatic indexical’) is its character plausibly viewed as a function from contexts to content. Even in cases of other guaranteed-to-refer indexicals such as ‘here’ and ‘now,’ it is clear that characters merely *constrain* content, and fall far short of *determining* content. (For example, by ‘here’ do you mean this room? This building? This city? This country? This planet?) This is even more evident with, say, a discourse-initial use of ‘She is German’ in a room in which there are many females (though the speaker has one specific person in mind, whether or not it is contextually obvious to the audience which person that is).

Meanings merely constrain extensions, they do not determine extensions. The possibility of mistaken associations on this point—when it comes to slogans along the lines of ‘characters determine contents,’ or ‘characters are functions from contexts to contents’—amount to another reason to be wary of applying those terms to sub-sentential expressions. (Cf. note 13 for arguments against indexicalism building from exactly this point.)

At the sentential level, though the distinction between sentential characters (i.e., context-independent, compositionally determined, sentence-sized meanings) and sentential contents (i.e., the specific truth-conditions expressed in context) is an indispensable tool which will be availed of at several junctures. The phenomena of indexicality (e.g., ‘Today is rainy,’ ‘She is German’) and

modality (e.g., ‘The current Prime Minister of Canada might not have been charming’) drive home the importance of this distinction. Crucially, for our purposes, it is at least an intelligible possibility that, for example, *a priori* is properly understood as a property of sentential characters while necessity is a property of sentential characters.¹⁴

[§]

In addition to the two semantic dimensions of meaning/extension and sentential characters/sentential contents, the externalist challenge necessitates distinguishing another semantic element—that is, the reference determiner. For the externalist challenge challenges precisely any traditional, Golden Age conception of the relations between meanings and extensions. It seems evident that, at least in certain cases, there are mechanisms of reference determination at work which bear little in the way of constitutive connections to the meaning (i.e., to that which needs to be grasped in order to be counted as competent with the relevant linguistic expressions).

The paradigm cases for motivating this distinction between meaning and reference determiner are natural kind terms (e.g., ‘aluminum,’ ‘elm’), particularly in their deferential uses by non-expert speakers. Here it seems that what counts as understanding the meaning falls well short of the ability to decisively delineate the extension. (You don’t need to be able to expertly distinguish elms or aluminum from similar things, in order to be counted as competent with the terms.) Similar points also seem to apply to proper names (e.g., ‘Aristotle,’ ‘Feynman’).

These considerations dovetail with certain lessons learned in the semantics of modality—specifically, when it comes to differences in extension, while meaning is held constant, across contexts of evaluation. Distinctions between meaning, extension, and reference determiner, and their impact on *a priori* and related notions, will pop up again in §3.2, and will be relevant to various points in Part III.

§2.4: INTERIM OVERVIEW

We have reached a plateau now, a base camp from which it is worthwhile to stop and look ahead at the climb ahead of us, reflecting on and integrating what we have taken on board in the course of our journey so far.

Our goal is to show that a constitutive *a priori* orientation can address the challenge of revisability and the externalist challenge, and yield a non-obscure but adequate and comprehensive epistemology. This is an understanding-based, rather than acquaintance-based, approach to the *a priori*; whether it should be thought of as a variety of moderate rationalism or of

moderate empiricism remains to be seen. Many pertinent details are yet to be worked out; the preliminary work executed in this present chapter will prove useful toward those ends.

Next, to draw out some ongoing morals. To ask whether something is necessary is to ask of a proposition whether it could be (or could have been) false, if contingent matters had gone otherwise. For example, provided that there could not possibly be fire without the presence of oxygen, then it is a necessary truth that fire requires oxygen. It is by now a familiar, Kripkean point that (given a minimal dose of realism) metaphysical necessity has nothing to do with anyone's knowledge of anything. Many consequences of this point will be further unpacked throughout; but now, to begin, here are two important ones.

One corollary worth underlining is that necessity may well be steadfastly indifferent to the challenge of revisability. That which is necessary is not revisable, for it is precisely the things which are most firmly bolted down that that concept is tailored to single out. (Of course, any particular agent's, or community's, guesses as to what is necessary may be revised over time, as a function of new evidence, insight, etc.; but that nowhere near entails that the truth-conditions of 'It is necessary that P' change over time. Necessity is the hardest of steel.) This is perhaps the central, deep important difference between necessity on the one hand and analyticity and *a priori* on the other hand—on the kind of understanding-based, constitutive *a priori* orientation developing here. Meanings can evolve, and what is taken to be self-evident varies widely across agents and communities; but metaphysical necessities are mind- and language-independent matters, steadfastly indifferent to what anyone thinks or says about them.

Secondly, and relatedly, metaphysical necessity is unlike the semantic and epistemic modalities in being in a certain sense vehicle-indifferent. That is, the semantic or epistemic modal status of a given proposition might be sensitive to certain features of the sentences which express it; but this seems to be decidedly less so for the question of necessity. (In other words, sentential contents matter for questions of necessity; and here sentential characters seem to be relatively unimportant. While, in contrast, not only sentential contents but also sentential characters are crucial to questions of analyticity and *a priori*.) Consider the following pairs:

- 1a. I am here now.
- 1b. Arthur is in the A.C. Hunter library at noon on October 18, 2017.
- 2a. Water contains hydrogen.
- 2b. H₂O contains hydrogen.
- 3a. All cordates are renates.
- 3b. All cordates are cordates.

It is controversial exactly how we ought to understand the precise content of the propositions expressed (in these cases as of many others), and it is beyond the scope of this present work to try to conclusively settle such debates within the philosophy of language. However, all that this present point requires is that it is possible that (at least some of) those pairs express the same sentential content but differ with respect to the questions of analyticity or *a priori*. Most importantly: given that they agree in truth-condition, then either member of each pair is necessary iff both are; but, on the constitutive *a priori* orientation, no analogous point holds of their status with respect to analyticity or *a priori*.

So, in this sense, necessity is, while analyticity and *a priori* are not, vehicle-indifferent. The above pairs of sentences are alike in metaphysical modal status. It is extensions, as opposed to meanings—that is, what the statement is about, as opposed to the way in which it characterizes or determines its extensions—which are crucial, when it comes to metaphysical modality. If a proposition is necessary, then it is so regardless of how it is characterized or expressed. In contrast, in the semantic and epistemic cases, the guises of different sentences can make a difference (even if they still might express the very same proposition).

[§]

Of our three focal notions, analyticity is the least vehicle-indifferent. Even the most slight and subtle difference of sentential character can be relevant here. In particular, in Part IV I will argue that [1a] is analytic while [1b] is not (even though both are true, and their content [in context] are at the very least closely related [if not identical]):

1a. I am here now.

1b. Arthur is in the A.C. Hunter library at noon on October 18, 2017.

As we will see, analogous claims could be defended for several other similar pairs.

To ask whether something is analytic is to ask: Does the meaning of the constituent bits, plus the mode of composition, suffice to ensure that the statement expresses a truth? Alternatively, is the denial of this contradictory? (Could a grandmother be childless? Or a bachelor married?) This is—in part, but centrally—a question about vehicles. One of the levels which it is so important to distinguish, when it comes to questions about analyticity (i.e., the conventional links between a sentence S and the proposition P it is used to express, on the one hand, and the truth-conditions of P, on the other hand) precisely concerns the linguistic vehicle of the expression.

Finally, to ask whether something is knowable *a priori* is to ask whether one is justified non-empirically in judging it to be true. First and foremost, *a*

priority is a property of the justification for a belief. Can this be known to be true without experience of what it is about? (For example, Could one steal one's own property?) Can this be known independently of any sensory experience? (For example, Could something think but not exist?)

Here the vehicle is not so clearly the issue, as in the case of analyticity, but it is still definitely relevant. [1a] (i.e., 'I am here now') is a relatively strong candidate for *a priori*, but it is far from obvious that [1b] (i.e., that I am at this—or indeed any—specific location at this—or any—specific time) should be so-classified. The epistemic status of a proposition seems to be relative to not just its intrinsic, truth-conditional content, but also to the guise or specific way in which the proposition is expressed. Hence, epistemic modality is also vehicle-relative; sentential characters clearly matter here too. (Relatedly, and as the externalist challenge drives home, as agents who fall decidedly short of omnipotence and infallibility, it is arguable that we could fail to recognize that what two distinct statements express is truth-conditionally identical. This will be further unpacked below in Parts III-IV; cf. especially §§6.4–5.)

[§]

So, then, on this framework-relative, understanding-based, constitutive *a priori* view, when it comes to both the issue of revisability, as well as the issue of vehicle-indifference, we have a clear split between metaphysical modality on the one hand and the semantic and epistemic cases on the other hand.

Whereas metaphysical modalities solely concerns extensions, truth-conditions, sentential contents, the semantic and epistemic cases also and essentially concern the dimension of meaning, characters, and the frameworks of meaning which are involved in understanding and communication. This is the root of some core differences, when it comes to revisability and vehicle-dependence, since our frameworks are themselves constantly subject to revision in light of new evidence and unforeseen connections, and concepts themselves (in addition to linguistic conventions pairing them with sounds and/or marks) are subject to a certain distinctive sort of evolution. (For example, what humans have believed about aluminum and metals has changed over time, and it is plausible that the meanings of the terms 'aluminum' and 'metal' have also evolved, but aluminum itself has remained constant and unconcerned throughout this ongoing process. If it is necessary that aluminum is a metal, this did not become the case as a result of any discovery or conceptual evolution.)

A further important aspect of this situation, also to be developed in depth below, concerns the constitutive ties between the semantic and epistemic cases—that is, between analyticity and *a priori*. One way into these ties is to consider the relation between what we have been calling 'semantic intuition' and 'rational intuition.' Though both alleged phenomena are controversial,

and both terms are variously employed, there is some reason to think that the terms ultimately target the same underlying thing: that is, understanding of concepts grounding the justification of beliefs. At a minimum, more conservatively, one might hold that semantic intuition is a sub-part of rational intuition; semantic competence is always and essentially a part of *a priori* justification. Ultimately, our beliefs are constituted by our meanings. The root of the close links between analyticity and *a priori* lies in the fact that the (epistemic) frameworks which are constitutive of our beliefs are themselves constituted by (semantic) meanings.

These deep, constitutive links between semantic and epistemic immunity to counterexample will amount to some strong reasons to think that all analytic truths are knowable *a priori* (though not necessarily the converse). Prior to explicitly addressing such questions in Part IV, much more ground needs to be excavated, when it comes to developing the relevant notion of a framework, and related conceptions of understanding-based accounts of the semantic and epistemic modalities.

NOTES

1. See chapters 5 & 6 of Grayling (1997) for extensive discussion of the comparative strengths and weaknesses of various theories of truth. I am not claiming that there are no outstanding technical questions for deflationists (e.g., semantic paradoxes), but they are outside the scope of this present work.

2. There is a related discussion of what Coffa (1991: 9) calls the ‘chemical theory’ of concepts, as it pertains to analytic truth, below in §3.2.

3. Cf., for example, Soames (2002: Ch’s 9–11) for extensive discussion.

4. There is discussion of this distinction in §3.1, and it crops up again in §7.3.

5. Of course, it is perfectly intelligible, and sometimes quite significant, to counterfactually vary linguistic conventions—for example, suppose ‘arthritis’ meant something distinct from what it actually does (cf. Burge [1979]). The present point is just that this latter sort of thought experiment is quite different from the kind of modal inquiry in which our interest is in the expression’s extension—for example, might water have turned out to be composed of XYZ? In these latter cases, it is crucial to hold fixed the meaning of ‘water’.

6. Boghossian calls this a metaphysical conception of analyticity, as opposed to a distinct epistemic conception. There is much more on this issue in §3.2.

7. The exception to this is indexicality (i.e., words like ‘this’ or ‘today’, whose extension changes from use to use, while their meaning nonetheless stays constant), which will be discussed in some depth immediately below.

8. Cf. Loewer (1997) for a good general accounts of this dialectical situation.

9. There are various other important complications to the Golden Age conception of the relation between context-independent meanings and propositions expressed

in context—such as context-sensitivity and semantic underdetermination. Indexicality and modality are the most important for our purposes, for reasons to be detailed below, and in any case are sufficient to develop the main departures from the Golden Age. For more on context-sensitivity, cf. Searle (1978), Recanati (2004), and for a discussion of semantic underdetermination cf. Bach (2005).

10. Cf. Stanley (2007) for development of a view generally classified as ‘indexicalism’, which takes indexicality to be a pervasive norm and not a circumscribed, distinctive sub-case.

11. Here see Kaplan (1989: 539): ‘The bearers of logical truth and of contingency are different entities. It is the *character* ... that is logically true [in cases such as “I am here now”], producing a true content in every context. But it is the *content* ... that is contingent or necessary.’

12. See Sullivan (2015a) for an overview.

13. Here I am in broad agreement with Schiffer (2003), Bach (2005), Neale (2007).

14. Cf. note 11, or consider Donnellan’s (1983) talk of the *a priori/a posteriori* distinction only applying to a proposition through the guise of the various sentences which can (in context) express it.

Part II

**KEY VARIETIES OF IMMUNITY
TO COUNTEREXAMPLE**

Chapter 3

Necessity and Analyticity

The next two chapters are dedicated to further analyses of the substance of the concepts necessary, analytic, and *a priori*. I will also consider some arguments for and against the intelligibility and worth of each concept. Looking ahead, one aim of these two chapters is to pare down these notions, to isolate their core out from amongst many closely nested questions and issues. The general goal is to arrive at concepts that are refined enough so that specific theses about their substance and interrelations can be tested, but yet not so over-refined as to be out of touch with their broad historical roots.

§3.1: MORE ON NECESSITY

‘Necessary’ is opposed to ‘contingent’ or ‘accidental.’ Something is necessarily so if and only if it could not possibly not be. Necessary truths are absolutely firm, unalterable; they are fashioned from the hardest of steel. They just simply are the case, regardless of time or place. They could not be otherwise, irrespective of how contingent matters of fact might be altered. We have no choice or influence in the matter.

Two concepts that are closely tied up with necessity, both historically and conceptually, are eternal truth and universal truth. A link between necessary and eternal is evident in Plato’s works, for instance—he talks of necessary truths as timeless, as not subject to temporal change. To cite another example, Leibniz explicitly uses ‘eternal truth’ and ‘necessary truth’ as synonyms (cf. Pap [1958: Ch. 1]). ‘Eternal truth’ contrasts most specifically with ‘historical truth’; a historical truth is one whose truth-value is contingent upon specific historical developments, and so could vary across the ages. So there is reason to think that a truth must be eternal, not subject to change over time, to be necessary.

It is also commonly claimed that all necessary truths are universal. The idea here is that, if a truth just applies to some particular matters of fact, at a certain time or place, or in virtue of certain specific and local contingencies, then it is not necessarily the case. To illustrate, suppose it is claimed that it is necessary that the conjunction of three conditions—let us call them A, B, and C—precipitate the occurrence of some further condition D. (D could be, for instance, a specific disease, an economic recession, or a forest fire.) The idea that all necessary truths are universal comes down to this: the claim of necessity is that, absolutely anywhere that A plus B plus C occurs, D occurs. If the claim is qualified, so that A plus B plus C lead to the occurrence of D only at a certain place or time, or only given certain other specific local circumstances, then, given that universality is a criterion of necessity, such qualifications undermine the claim to necessity.

To be sure, though, claims of necessity do not depend on there actually being any specific minimal number of instances. (Analogously, the claim ‘Trespassers will be prosecuted’ is not rendered *false* or *meaningless* if no one, in fact, ever happens to trespass.) A plus B plus C might only actually co-occur once, or might not co-occur at all, without undermining the claim that the asserted connection to D necessarily holds. So, for instance, the claim that the presence of oxygen is necessary for fire is not a claim about how many fires there are; and, if now true, it did not just become true the first or tenth or millionth time there was a fire. Rather, it was true all along, determined by the very nature of the phenomena.

One factor that complicates this link between necessity and universality we might call Heraclitus’ worry. Heraclitus is said to have remarked that one can never step into the same river twice. One thing he is taken to have meant by this remark is that, strictly speaking, no situation ever occurs more than once; if this is so then all phenomena are specific and local, and the very idea of universal truth might seem to be useless artificial abstraction. For any event in the real world—diseases, recessions, fires, and so on—there are an awful lot more than three causally relevant factors to consider. It may well be the case that, for any particular phenomena, the causally relevant factors are never exactly repeated. If so, all connections between phenomena would be, in a sense, contingent on local circumstances, and universality would seem idle and insignificant.¹

Note, though, that although these considerations serve to underline the complexity of the study of metaphysical modality, and to guard against oversimplification, they do not undermine the interest or worth of the core notion of necessity. These complexities are real and important, but they do not entail that there is no such thing as factors that necessitate diseases, recessions, or fires, or that it is folly to keep looking for them.

The notion of generality lies at the base of these links between necessary, eternal, and universal. It is widely held to be part of the content of the concept

of necessity that, in principle, necessary truths apply to more than one particular situation or individual. The scope of a claim of necessity includes all places and times; following Plato, this is sometimes glossed as the thought that necessary truths are outside the bounds of space and time. And so, the more instances we observe of a connection between two things—oxygen and fire, say, or cessation of heartbeat and death—the stronger is the case for the claim that the connection between these phenomena is necessary, not just accidental.

Necessity, *per se*, has nothing to do with anyone's knowledge of anything, or with what any linguistic expressions mean. It has to do with the nature of things, and while we certainly try to fashion our beliefs to track the mind-independent nature of things, and to craft our terms so that we can accurately discuss the phenomena which matter to us, that something is necessary is not a claim about knowledge or meaning, and does not immediately entail any such epistemological or semantic claim. It is a matter of metaphysics, and is something about which all humans could be mistaken or ignorant, or lack the conceptual means to engage with in their thought and talk.

[§]

The distinction between essence and accident is one of the central distinctions in modal metaphysics. Take any individual, or kind, or phenomena. What changes can it persist through, and what changes would be so drastic as to preclude the sustenance of its identity? Sure, Gillian might have been a lawyer instead of a teacher, but presumably she couldn't have been a rock; sure, the Olympic Games would still be the Olympic Games if cross-country wrestling were admitted as a sport, but probably not if it no longer contained any sports and did not involve international competition. Similarly for diseases and sub-atomic particles and political parties: an integral part of understanding any kind of phenomena is to sort what is essential to it from what is accidental.

There is much controversy concerning the precise content and intelligibility of these kinds of modal attributions, and the relations between them. These and related questions form some ancient threads in metaphysics, still vibrant today. For instance, consider the question: Do individuals have an essence? Aristotle held that essences are species-wide properties. (Indeed, 'species' and 'essence' come from the same root word, the former stemming from a Medieval Latin variant on the older Greek term.) All tigers, and all eyes, share an essence, down this avenue; the essence of something is what it does; and all tigers, all eyes—all members of any one species—should in this sense be seen as *doing* the same thing. Therefore individuals do not really have distinct essences, on Aristotle's view. Kripke (1972), in contrast, defends a very different view, according to which individual humans (and some other kinds

of objects) do have unique individual essences. There is something unique about being me, or being a particular tiger, on Kripke's view. He makes some substantial and controversial conjectures about individual essence.²

Those who believe in distinct individual essences are called 'haecceitists.' (The word comes from the Medieval Latin translation for the Ancient Greek word for 'this.')

Haecceitists believe that there is something that constitutes the essence of a particular thing, something it necessarily has, without which it would not be itself, and that nothing else could have. There is something unique and essential about an entity, which endures through the change of all accidental properties. Anti-haecceitists reject this claim. Anti-haecceitists are skeptical of the coherence or usefulness of the distinction between individual essence and accidental properties.³

[§]

I will next discuss the way in which the exact strength or force of modal terms like 'necessary' can vary with context. In most cases, modal terms are not used in a completely unrestricted manner, but are rather implicitly restricted to some (more or less vaguely defined) contextually salient set of possibilities. For example, suppose I sincerely utter:

1. I would really like to make this \$50,000 investment, but I am afraid that it is just not possible right now.

Clearly, not the whole of the vast expanse of metaphysical possibilities is relevant to the content of such a statement. Rather, only those possibilities in which the constraints on my resources are (more or less, roughly) held constant are relevant. It would be obtuse for one to respond that [1] is false because there is a perfectly possible situation in which I could easily make this investment (e.g., if I had won the lottery yesterday). Such remote possibilities are irrelevant to what statements like [1] express, which is to say that statements like [1] involve implicit restriction to a contextually salient proper subset of the set of all possibilities.

Another way to illustrate this phenomenon is with reference to counterfactual conditionals—that is, hypothetical 'if'/'then' statements whose 'if' clause is false. Counterfactuals, like probabilistic reasoning, involve reasoning about non-actual possibilities; and, from scientists to Monday morning quarterbacks, we use them all the time in everyday reasoning:

2. If Hitler had not been born, then World War II would not have happened.
3. If the Seahawks had punted instead of gambled on the last play of the third quarter, then they would have won the game.

Counterfactuals also clearly instance this phenomenon of implicit restrictions on possibilities. [2], for example, asserts that Hitler was the, or at least a central, cause of World War II. Again, it would be obtuse to argue against [2] on the grounds that even if there were no Hitler, still it is possible that World War II was started by, say, a one-armed Australian sheep farmer. [2] is about the loosely defined set of possibilities that are most like the actual world except that Hitler does not exist; it is not the claim that it is metaphysically impossible for there to be World War II but no Hitler. Similarly for [3], which does not assert that it would contravene the ultimate laws of the universe for the Seahawks to both punt and lose. Only a selected, target set of possibilities are involved, in the context of such discussions.

This is an important point to beware of, generally. In ordinary discourse, there is much vagueness, ambiguity, and shiftiness inherent in metaphysical modal claims. When one encounters such a claim, one needs to be keen to various contextual clues, in order to determine precisely how restricted, or unrestricted, the scope of the claim is intended to be. The cases of analyticity and *a priority* differ fairly drastically, in this respect. While, as we will see, they do instance a distinctive and pervasive sort of framework-relativity, they are not subject to exactly the same sort of context-dependent shiftiness, illustrated by [1]–[3] above.

[§]

A related refinement is that there are different kinds, or strengths, of necessity. One distinction which is drawn in this terrain is that between physical and metaphysical necessity. ‘X is physically necessary’ means that X is necessitated by our actual laws of nature (such as $E = MC^2$, or $G = M_1 \times M_2 / D^2$). ‘Y is metaphysically necessary’ means that Y is necessitated by the laws of metaphysics (such as that everything is self-identical, or that no individual simultaneously both has and lacks a given property). Physical necessities hold throughout the set of physically possible worlds, that is, the worlds in which the physical laws of the actual world hold. Metaphysical necessities hold throughout all metaphysically possible worlds. Reichenbach (1953) gives a nice way to illustrate this distinction. Consider:

4. There is a solid one-ton sphere of gold.
5. There is a solid one-ton sphere of uranium 238.

Both are actually false. However, [5] is physically possible, while [6] is not. That is, while the sphere of gold could exist—maybe Oprah could build one, if she wanted to—the sphere of uranium could not, as a matter of physical law, because it is radioactive, and would explode. However, if one thinks that

the uranium sphere could exist, in possible worlds in which the laws of nature are different, then one thinks that [6] is metaphysically possible, even though physically impossible.

Intuitively, the idea behind this distinction between physical and metaphysical possibility is that there is some degree of contingency to the actual laws of nature. It is in some sense possible that, say, gravity or magnetism might have worked differently, or might not have been at all, and that would not spell the end of the world. If so, we are describing physically impossible but metaphysically possible worlds. Or, again, take Planck's constant, which gives the ratio of the frequency of radiation to its quanta of energy. Planck's constant has an approximate value of 6.625×10^{-27} . Now surely, one might think, things did not have to be this way. Planck's constant might have had the approximate value of 6.624×10^{-27} , or 6.626×10^{-27} , and the universe would still exist (in more or less the same fashion). If that is so, then there are metaphysically possible worlds in which Planck's constant has a slightly different value; but these would be physically impossible worlds, because the actual laws of physics would not hold. If pigs can fly, and cows can jump over the moon, then this happens in physically impossible worlds. (In the actual state of things, pigs have not shown any aeronautical capability, and cows cannot jump over squat.)

[§]

Metaphysical modality was an area of vibrant debates in the Ancient and Medieval periods in Western philosophy, but was to subsequently endure various degrees of neglect or hostility, until well into the twentieth century. For example, during the Modern period, epistemological questions were given foundational priority over metaphysical ones; and, during the first half of the twentieth century there was active antipathy toward many such metaphysical notions—some of the reasons for which will be discussed below in §3.3.

The notion of mind- and language-independent necessity became much more widely accepted in the latter decades of the twentieth century, after enduring this period of relative neglect. Developments pertaining to the sophistication of our understanding of semantic matters makes it possible to detect fallacies and confusions in several varieties of argument against the coherence of metaphysical necessity (cf. especially §3.3 below), and rigorous semantics for modal discourse were developed.⁴ We post-Kripkeans are more at home with many varieties of modal attributions; even if there is still much controversy as to understanding the exact content and import of such attributions. That is, metaphysicians of many different stripes and orientations can agree that it is not possible for something to have the property 'square' but not be four-sided, or to have the property 'water' but not be H₂O. It is at least

commonplace, if not orthodox, in these liberal times, to accept mind- and language-independent necessary connections among properties.

§3.2: ANALYTICITY REVISITED

A truth is analytic if and only if it is true by virtue of the meanings of the terms involved; alternatively, the denial of an analytic truth is contradictory. Analyticity is a kind of semantic guarantee of immunity to counterexample. Candidates for analyticity include:

1. Squares have four equal sides.
2. No grandmother is childless.
3. Bachelors are unmarried men.

Recall Plato's problem: Our experience is particular and limited, but yet we still manage to attain knowledge of some universal, necessary truths. As mentioned in chapter 1, many philosophers—including especially empiricists who countenance universal, necessary knowledge—have hoped that the notion of analytic truth can ground a non-obscure but adequate solution to Plato's problem. The thought is that, in certain privileged cases, just grasp of meaning is sufficient to justify the belief that the statement is universally and necessarily true. (In other words, semantic intuition is not an obscure, non-empirical faculty of mind, and it might be capable of delivering all that was wanted of rational intuition.)

The term 'analytic' comes from the root 'analysis,' another word with Ancient Greek roots. An analytic approach to a problem proceeds via separating it into component parts or constituent elements. 'Analytic' contrasts with 'synthetic.' In the case of a synthetic statement (such as that Neptune has four moons or that Alice is a grandmother), more than just analysis of the constituent terms is required in order to judge whether it is true. 'Synthetic,' from 'synthesis,' means having been put together. Analysis is a process of taking something apart; synthesis is a process of putting something together.

What exactly, is analyticity a property of? Sentential characters are not the best candidates for analyticity, because sentences per se are not true or false (a fortiori not true or false in virtue of anything). Rather, a use of a sentence in a context can express something, and the question of truth or falsity arises with respect to what is expressed. So, is analyticity then a property of what is expressed—that is, a proposition, or sentential content? First and foremost, my reason for avoiding that option is the point from §2.4 about vehicle-indifference (i.e., relevantly different sentences can be used [in context] to express

the same proposition). For a variety of sorts of cases, some of which will be important below, it is important to leave open the possibility that different ways of expressing the same proposition might differ with respect to either or both semantic or epistemic status. Again, examples include:

‘Cordates are cordates’ versus ‘Cordates are renates’.

‘I am here now’ versus ‘Arthur is in St. John’s on October 18, 2017’.

‘H₂O contains hydrogen’ versus ‘Water contains hydrogen’.

Even if such pairs are taken to express the same proposition, still there might be good reasons to count *only* the first member of each pair as analytic (and/or *a priori*).

So, the question of truth brings propositions into the picture; for this question does not arise at the level of sentences. However, given that there can be relevant differences between distinct sentences which express the same proposition, then propositions are not the only thing in the picture. Thus, the bearer of analyticity will have to encompass both a proposition (in order to accommodate the question of truth) and the means of expressing that proposition (in order to accommodate the non-truth-conditional element in analyticity). That is, analyticity is not exclusively a question about sentential characters or about sentential contents, but essentially includes both dimensions.

Hence, I take the primary bearers of analyticity to be statements, where a statement is a particular, dated use of a sentence to express a proposition. Analyticity, then, is a property not solely of symbols, or of contents, but of specific uses of symbols to express contents.⁵

[§]

At this juncture, I should explicitly flag two fundamental distinctive properties of analyticity, as distinct from necessity or *a priori*. Both will be revisited at several junctures below. The first pertains squarely to Plato’s problem. I will call it the ‘from coherence to worth’ worry: that is, for the case of analyticity, more so than for the cases of necessity or *a priori*, there prevails the worry that even if the notion is coherent, and uncontroversial examples could be found, still it is not yet clear that analytic truth is of any philosophical interest, as opposed to being merely verbal trifle. One would hardly hear either ‘Sure, that is a necessary truth; but necessity itself is of no philosophical interest,’ or ‘OK, fine, that is *a priori* knowledge; but that category of knowledge is insignificant.’⁶ However, this kind of worry is commonly pressed for the case of analyticity: ‘Sure, OK, I grant that “All bachelors are married” is immune to counterexample in virtue of meaning. But how does that get us even one iota toward solving Plato’s problem?’ Again, this ‘from coherence to worth’ worry

threads throughout several strands of the ensuing discussion. One reason why I flag it here is that it is crucial—for both proponents and opponents—to distinguish objections to the coherence (or intelligibility) of analytic truth from objections to the worth (or usefulness) of the notion.

A second distinctive complication is that analytic truths are essentially about meanings, and meanings are notoriously difficult things to talk about in a non-contentious, non-tendentious way. Not only have there been considerable and influential skeptics about the notion of meaning, there have also been lots of non-trivial differences between the non-skeptics, as to how to understand this complex notion. So, it seems that lots of honing and framing of the notion of meaning is a precondition for a worthwhile discussion of analyticity—again, as opposed to the cases of necessity and *a priori*. Of course my point is not that there are no substantive philosophical questions about the notions of metaphysical reality or of epistemic status, but rather that the notion of meaning is relatively worse off, when it comes to getting a pre-theoretical handle on the very idea, in order to get a philosophical discussion rolling.

So, one of the reasons why there was relatively little discussion of the notion of analytic truth in the latter half of the twentieth century is that there is such a wide variety of distinct theoretical approaches to meaning. There are philosophers who take the basis of meaning to be sense, reference, use, intentions, truth-conditions, causation, teleology, etc., and each of these orientations has distinct sub-varieties. As a result, it is difficult to be precise about exactly what ‘truth in virtue of meaning’ comes to without assuming a specific and controversial take on the nature of meaning. To proceed in terms of any of sense, reference, use, intention, truth-conditions, causation, teleology, etc., would be controversial; while to proceed in terms of all of them would be untidy and vague.

For present purposes, though, it is neither necessary nor desirable to pin down ‘meaning’ more precisely. My view is that if meaning is determinate, then there will inevitably be a class of statements that is true in virtue of meaning.⁷ (There may be borderline cases, of course, but the present point is that there will still be paradigm cases. From the fact that there is a broad spectrum of shades of grey, it hardly follows that nothing is either black or white.) That is, whichever approach to meaning to which one subscribes, as long as it is possible, if not common, for some specific content to be semantically expressed with a statement, then the brute datum described above (of immunity to counterexample in virtue of meaning) is bound to arise. Hence, it is appropriate for this discussion to stay as ecumenical as possible about the correct theoretical account of meaning.

Of course, some of the most influential critics of the notion of analyticity have also (and not coincidentally) been skeptics about the determinacy of meaning. (Quine is a seminal case in point.) I will not say much that is directly addressed to refuting such meaning-skepticism here, but will rather

confine myself to two pertinent points. One minimal point is that it is important to distinguish the following two questions, with respect to any use of a sentence (in context):

1. Can we know for sure exactly what a certain speaker is intending to communicate, with a certain utterance?
2. Is there a determinate meaning semantically expressed?

The first (interpretive) question is a complex and difficult knot, which it would take considerable hard work to even begin to untangle. However, even if the interpretive question should ultimately be answered in the negative, that does not mean that the second (semantic) question should also be answered in the negative. There are lots of familiar, intuitive reasons to think that the second question should be answered in the affirmative,⁸ and that is enough to get the problematic of truth in virtue of meaning up and running. Again, if meaning is determinate, then there will inevitably be a class of statements that is true in virtue of meaning. (With a nod back to §1.3, this is to distinguish between epistemic from metaphysical questions, in this particular domain.)

Second, also note that while meaning-skepticism was an option for a naturalistically inclined philosopher in the mid-twentieth century, when behaviorists and structuralists still ruled the human sciences, this is no longer so at this juncture, since the cognitive revolution. The reason is that what are currently the most successful research programs in the study of cognition are up to their minds/brains in intentional semantic notions. To dismiss these research programs as wrong-headed is to presume a perspective outside of science from which science can be evaluated. So, a naturalist must respect these research programs, and to do that is to reject meaning-skepticism.⁹

Of course all that still leaves us with the hard work of explaining how something which might seem as trivial as ‘truth in virtue of meaning’ has any promise to afford a substantive and weighty answer to Plato’s problem. A successful defense of the coherence of analytic truth would not yet amount to a positive case in favor of its significance or worth.

[§]

While the explicit employment of the notion of analyticity is a relatively modern phenomenon, the core idea underlying the notion has been around for a long time. For one thing, many propositions that have been of perennial interest in philosophy are candidates for analyticity, such as:

1. No proposition is simultaneously both true and not true.
2. Every even number is divisible by two.
3. Murder is wrong.

The strategy of using something like analytic truth to fashion an answer to Plato's problem is almost explicitly formed in the work of some Modern empiricists, such as Hobbes and Hume. Hobbes seems to have espoused the view that the root of all necessity lay in language, that what we think of as necessity is properly traced back to linguistic conventions.¹⁰ Variations on this conventionalist theme have been players on the scene ever since.

With Hume, we come very close to an explicit statement of this style of answer to Plato's problem:

All of the objects of human reason or enquiry may be naturally divided into two kinds, to wit, *Relations of Ideas*, and *Matters of Fact*. Of the first kind are the sciences of Geometry, Algebra, and Arithmetic: and, in short, every affirmation which is either intuitively or demonstratively certain. ... Propositions of this kind are discoverable by the mere operation of thought, without dependence on what is anywhere existent in the universe. ... Matters of fact ... are not ascertained in the same manner; nor is our evidence of their truth, however great, of a like nature with the foregoing. The contrary of every matter of fact is still possible. (1748: Sect. IV, Part 1)

This is a seminal statement of what many empiricists want to say about our universal, necessary knowledge: We have *a priori* knowledge of certain things because they are analytic—that is, because of intrinsic relations among the relevant concepts.

To be sure, though, what Hume calls 'relations of ideas' involves an undifferentiated mush of what I want to distinguish as necessity, analyticity, and *a priority*. Further, as Kant stresses, it is far from clear that such Relations of Ideas can afford the basis for substantive, informative, judgments (as opposed to merely *a priori* trivialities).

[§]

As with the cases of *a priority*, a defensible notion of analytic truth must steer clear of some of the features that have been historically associated with the concept. It is crucial to sharpen the concept of analyticity, because, even though the notion has a respectable and important place in the philosophers' toolkit, there are serious problems with many of its traditional associations. Thus, I must disavow some of the prevalent associations of 'analyticity,' and develop some refinements. While many would hold that it is precisely those associations which constitute the main philosophical interest in analyticity, I will argue that the conception which remains is not only defensible but quite substantive and worthwhile.

To begin this pruning process, I will start from a quote from Coffa (1991: 9–10):

One of the many ways philosophers have tried to understand meaning might be called the ‘chemical theory of [propositions],’ using an analogy occasionally found in the writings of Locke ... and Kant. According to this theory, [propositions], like chemical compounds, are usually complexes of [meanings], which may themselves be complex. ... Analysis is the process through which we identify the constituent [meanings which make up a proposition, and then, in turn, break down these meanings into their constituents]. It is a process that must come to an end ... in the identification of simple constituents. ... To know a concept fully, for example, is to define it; and definition is no more and no less than exhaustive and complete analysis.¹¹

I will not attempt to document the historical prevalence of this chemical theory. Rather, first I will briefly explain why, given the chemical theory of propositions, the analytic/synthetic distinction is both sharp and important. Second, I will point out that there are lots of reasons to reject the chemical theory. Finally, while this does suggest that the analytic/synthetic distinction is less sharp than many had thought, it does not entail that the distinction lacks importance for philosophy. Rather, the shortcomings of the chemical theory point to some valuable refinements to the notion of analyticity.

So, first: if the chemical theory of propositions were an accurate picture of semantic content, then the analytic/synthetic distinction would be crystal clear, and its philosophical importance would be absolutely paramount. The guiding idea here is that meanings, and the propositions which they compose, are objective, mind- and language-independent entities (with objectively identifiable compositional structures). Whether one is more inclined toward a dialogical sort of view that getting in touch with these objective entities is essentially an intersubjective affair, or toward a more subjectivist view that one isolated mind is capable of getting acquainted with them, either way there is an objective fact of the matter as to whether a statement is true in virtue of meaning (or, alternatively, of whether its denial is contradictory). The question of whether a certain judgment is analytic is, like the question of whether a certain sample of ore contains any iron, a matter for objective discovery.

Further, given this chemical theory of propositions, the analytic/synthetic distinction is rather vital to the enterprise of philosophical inquiry as a whole. Indeed, there is a strand within philosophy which takes the analytic/synthetic distinction to be definitive of the very subject matter of philosophical inquiry, as distinct from scientific inquiry. While both sorts of inquiry might be objectively factual, synthetic matters of fact are appropriately left to scientists, while the distinctive subject matter of philosophy is conceptual analysis. The unique expertise of the philosopher, on this orientation, lies in mining, amplifying, and tracing inferential patterns among analytic judgments, and in cultivating this kind of specific knowledge into general wisdom.

However, throughout the twentieth century, in various ways, it has become generally acknowledged that concepts and propositions are not so simple as the chemical theory would have it. To the contrary, languages are organic entities, and the conventions pairing expressions with meanings evolve over time. Languages, meanings, and conventions are non static, objective entities inhabiting some realm out there, waiting for us to get acquainted to them, but rather, are constantly being revised, expanding while here and decaying there. New meanings and conventions are ever suddenly being coined, and gradually dying off. These considerations greatly complicate any links between linguistic conventions and necessary truth.

Relatedly, the idea that there are objective mind- and language-independent facts about meanings and propositions is no longer generally conceded as unproblematic. Many philosophers hold that not just linguistic conventions, but meanings and propositions themselves, are, in some significant senses, human-dependent. To the extent to which the very content of the concepts at issue is, at least in part, up to us—and so can vary from time to time and place to place—the very idea of objective facts about analyticity is deeply problematic. The idea that there are timeless objective facts about analyticity now seems rather archaic, given the general recognition that languages and conceptual schemes are organic entities which evolve over time. (There will be much more discussion of this kind of conceptual evolution in Part III; especially §6.3.)

I will close this particular thread of the discussion by entering two critical remarks, pertaining to aspects of Kant's views which must be either rejected or substantially refined. The first point is that Kant was an adherent of what was above called 'the chemical theory of concepts,' as is evidenced by another definition of analyticity he endorses which appeals to the metaphor of containment—that is, a judgment is analytic iff the predicate concept is contained in the subject concept. Not only does this criterion just apply to a limited class of statements,¹² it also seems to clearly presuppose the picture of concepts as static entities out there waiting for someone to get acquainted with them, of there existing timeless objective facts about what is or is not analytic, which is now generally rejected.

So, there is a real definite lack of substantive criteria for analyticity here. Notoriously, controversies as to whether a contested candidate should be considered to be analytic or synthetic are not resolvable on the basis of anything in the Kantian corpus. For this reason, debates between Kantians and their opponents had, by the mid-twentieth century, 'settled, or bogged, down around a handful of particular cases: for example, "nothing can be red and green all over" and the transitivity of "earlier than"' (Quinton 1963: 31). Contrasting positions on the synthetic *a priori* status of such cases deteriorated

into articles of faith, because of the ineffectuality of the extant ways of drawing the analytic/synthetic distinction.

For these and other reasons, Kant's conception of the distinction has been subject to an awful lot of criticism.¹³ So, even though the view developed in Parts III and IV owes a lot to Kant, below I will depart from the specifics of his views at several junctures.

For now, one important moral is that the analytic/synthetic distinction is no longer taken to be as crisp as Hume or Kant, say, thought it was. Rather, it is perhaps more plausible to view it as a continuum with paradigm cases at the extremes and a range of shades of grey stretching in between. Even so, my view is that, suitably refined, it is still of philosophical importance.

For instance, despite the demise of some of these oversimplified ideas about meanings and propositions, there is still the brute datum—that is, we are justified in believing some statements that are universal in scope, and not subject to refutation by contingent happenstance, such as:

1. All squares have four sides.
2. No grandmother is childless.
3. Two is a factor of every even number.

These brute data provide strong *prima facie* paradigm cases of analytic truths. Among other things, this may yet hold the key to a satisfactory solution to Plato's problem. The idea is that grasp of the meanings of the constituent terms is sufficient to justify the belief in the truth of such statements. In such cases, understanding of meaning grounds the recognition of truth.

[§]

A couple of final refinements to analyticity, before turning to some claims staked by skeptics. One stone which has remained as yet unturned concerns how 'ampliative' an analytic truth can be, with Kant (1781) and Frege (1884) at the opposed poles. For Kant, analytic judgments cannot afford new knowledge, but rather just merely unpack the content of what is already known. In contrast, Frege derides 'the widespread contempt for analytic judgements' and its attendant 'legend of the sterility of pure logic' (1884: 24). Frege's metaphor that the fruits of analysis 'are contained in the definitions, but as plants are contained in their seed, not as beams are contained in a house' (1884: 101) illustrates this nicely. Logical and semantic analysis can afford new knowledge, as opposed to telling us what we already knew—as, indeed, anyone who has had to work hard on finding a derivation of a self-evident theorem of logic, or found a simple, elegant proof of something entirely non-trivial, can appreciate. So, Frege clearly rejects conceptions of analyticity which would tie it to lack of substance, to emptiness of content. Analytic

judgments can be ampliative for Frege (i.e., substantive increases in knowledge), unlike for Kant.

This is especially relevant to our ongoing ‘from coherence to worth’ struggles, for Frege (and many of those influenced by him) rejects the idea that analytic truths are trivial and uninteresting. Relatedly, it is also pertinent to ongoing questions about the ampliative power of semantic intuition itself, and ultimately to upon which side of the rationalism/empiricism divide a variety of the constitutive *a priori* orientation is situated.

Another barely yet turned stone: In §2.3 I distinguished the notion of ‘reference determiner’ from both meaning and extension; this notion has become important in the wake of the externalist challenge. In penetrating recent work, Russell (2008, 2012) has argued that three more precise notions should now replace the old ‘truth in virtue of meaning’—namely, (i) truth in virtue of character, (ii) truth in virtue of content, and (iii) truth in virtue of reference determiner. On her view, analytic truth is worth saving, and (iii) is the best way to do so, in the wake of reviseability, externalism, and other considerations.

In the context of the present inquiry, a main part of which is tracing the overlaps and distinctions between analyticity and the closely related concepts of necessity and *a priority*, I have opted for keeping meaning in the picture. Continuity with that aspect of the traditional discussion is both intrinsically desirable and extrinsically useful. Furthermore, the response to the challenge of revisability which I develop in Part III is rather different from Russell’s, one result of which is that the notion of a reference determiner (as distinct from meaning and extension) will play a more limited role in my work than in hers. This is hardly a criticism of Russell’s work,¹⁴ but rather just a note that, while I find her work very interesting and worthwhile, I will not be following it very closely, in this respect.

§3.3: SKEPTICISM ABOUT NECESSITY AND ANALYTICITY

Skepticism about metaphysical necessity is not hard to motivate, or to sympathize with. Claims about essences are relatively obscure and mysterious; and they are awfully difficult to conclusively support. They are rather remote from everyday experience and practical concerns. For these and other reasons, there have been skeptics about necessity for as long as there has been modal metaphysics. Next I will discuss two variants or aspects of this skepticism: first, the epistemological charge that claims of necessity could never be justified; and second, the stronger charge, partly semantic and partly metaphysical, that such distinctions as necessary/contingent or essence/accident

are incoherent. The first is easier to establish but less consequential, and the second would be quite serious, but is rather hard to establish.

[§]

There is some reason to think that we could never conclusively verify the claim that a proposition is necessary. By definition, a claim of necessity extends beyond the range of not only any particular individual's observation, but, even further, beyond the range of the observations of all actual individuals, past and future. Consider, for instance:

1. It is necessary that the amount of force an object will exert on another will rise in proportion to an increase in its velocity (i.e., the faster it's moving, the harder it'll hit).

[1] is not just a claim about all cases that have been, or will be, observed. It is the claim that it is a fact determined by the very nature of things and unconstrained by specific spatio-temporal contingencies, that it is not possible to increase an object's velocity without increasing its potential energy. This claim extends indefinitely through time, and throughout all possible ways the world could have been—that is, the scope of [1] is so broad that it applies regardless how the world might have been contingently different, wherever and whenever we choose to consider. Given that we can never conclusively verify claims of necessity, why should we believe in them, and what use are they?

Even if, strictly speaking, we can never conclusively evaluate such claims to necessity, a reasonable response is to admit that our judgments of necessity are fallible, and subject to change in the face of counter-evidence, as opposed to outlawing the term 'necessity' from rigorous and systematic thought. The more instances that are observed of a connection between two things—such as velocity and force, or fire and heat, or cessation of heartbeat and death, and so on—the less plausible it is that the connection is accidental, and the greater the reason one has to believe that the connection between them is necessary, is rooted in the very essence of the phenomena. Further, the positing of necessary connections, where we have good reason to, has great explanatory value. In general, the more well-confirmed hypotheses of necessary connections we have, the more comprehensive is our understanding of how things work, and the greater is our ability to predict future conditions, and to manipulate conditions to our advantage. Indeed, this kind of non-accidental, mind-independent connection between phenomena is precisely the target at which scientific inquiry (as well as much non-scientific inquiry) aims.

So, we can have very good reason to believe in claims of necessity, asymptotically approaching but never reaching certainty, even though they

extend beyond the range of our experience. The usefulness of the notion can be demonstrated by its conceptual role at the foundations of the scientific enterprise, which has led to countless useful improvements in our ability to predict future occurrences, and to manipulate natural conditions to our own advantage.

Regardless, it is important to underline that claims of necessity are not claims about knowability, verifiability, or usefulness. Those are important and difficult issues, but reasons to be pessimistic about them do not undermine the intelligibility of the notion of necessity. Further, it is crucial to separate them out from necessity, in order to isolate the relevant notion, before we can productively address these questions about how claims to necessity are established or what use they might be.

[§]

Hume is a seminal source of skepticism about metaphysical necessity. He is often read as arguing that such concepts are things that minds impose on, rather than discover in, nature. Nowhere in our experience do we observe causation, law, or necessity; rather, these are just parts of our conceptual repertoire that are useful in helping to organize our beliefs and theories about nature. Necessity is no more out there in the world than are centers of gravity or the equator, according to Humeans. These are one and all human constructions that have no real, objective mind- or language-independent correlatives.¹⁵

It is not possible to decisively counter Hume's objections. We are not in the realm of conclusive proof, when it comes to the hypothesis that notions like law, causation, and necessity do in fact have real, objective mind- and language-independent correlatives. But who should have the burden of proof here: the skeptic or the realist?

Well, there is no shortage of reasons to think that modal concepts are rather important, useful, significant. We will soon turn to showing that certain influential arguments to the contrary can be countered, given due attention to fine distinctions within the philosophy of language, as well as to distinctions between the epistemological and the metaphysical aspects of an issue. For another thing, people do engage in modal speculation all the time, and it seems reasonably clear what they are talking about—that is, how familiar objects and individuals would have fared under different circumstances, or would have endured alterations to their accidental properties. Hume's point is a valuable one, exposing the lack of conclusive support for many of our most fundamental beliefs in epistemology and metaphysics; but it is a long way from this point to categorical skepticism about any and all claims concerning the metaphysical modalities. (It is

a long way from a ‘lack of conclusive support’ to a ‘conclusive lack of support’.)

[§]

Throughout I have been careful to distinguish between the issues of: (i) the *coherence* of analytic truth, and (ii) the *worth* of analytic truth. I will likewise arrange this section along these two dimensions, when it comes to engaging with skepticism about analyticity. Further, as for the charge that the notion of analytic truth is incoherent, I distinguish between what I will call ‘the weak coherence charge’ and ‘the strong coherence charge.’ The weak coherence charge is the idea that the analytic/synthetic distinction has never been drawn in a comprehensive and satisfactory way. The strong coherence charge is the much bolder—and, consequentially, more difficult to justify—notion that the very idea of the analytic/synthetic distinction is irredeemably confused.

It would be difficult to deny the weak coherence charge. While one can clearly see progress in the sophistication of discussions of the analytic/synthetic distinction—from Hobbes to Kant to Frege, and on into recent work by, say, Boghossian (1997) and Russell (2008)—still I cannot imagine any proponent of analyticity claiming that the matter has been conclusively laid to rest. Indeed, it is hard to see how there could be consensus concerning the analytic/synthetic distinction until there was consensus concerning the proper theoretical treatment of the notion of meaning, and it is safe to say that neither philosophy nor any of the cognitive sciences is anywhere near such a state. Among other complications, the analytic/synthetic distinction can only be as firm as meaning is determinate (here compare the discussion of the chemical theory of concepts in §3.2). While I see little force in sweeping, global skepticism as to the determinacy of meaning, given the reasonable view that meaning is thoroughly context-sensitive (i.e., that meanings are malleable, relative to, and as mandated by, the context of utterance), any formulation of the analytic/synthetic distinction will have to be programmatic, hedged, shift, and at least somewhat imprecise.

[§]

However, the strong coherence charge (i.e., that the very idea of the analytic/synthetic distinction is irredeemably confused) is another matter entirely. Could it be conclusively established? When we factor out Quine’s dated behavioristic skepticism about meaning, and the aspects of his arguments which apply not intrinsically to analyticity per se but extrinsically to some of the things to which certain empiricists attempted to apply this notion, what arguments are there for the strong coherence charge?

It must be noted that the degree of even Quine’s commitment to the strong coherence argument is questionable. His allegiance to at least the weak coherence point is firmly established in his (1951: 32), and is nicely stated thus:

I do not know whether the statement ‘Everything green is extended’ is analytic. Now does my indecision over this example really betray an incomplete understanding, an incomplete grasp of the ‘meanings’, of ‘green’ and ‘extended’? I think not. The trouble is not with ‘green’ or ‘extended,’ but with ‘analytic.’

However, there are passages in which it sounds like Quine (1951: 41) endorses something much bolder:

That there is such a distinction to be drawn at all is an unempirical dogma of empiricists, a metaphysical article of faith.

As Grice & Strawson (1956) quickly pointed out, though, the argument in Quine (1951: Parts I–IV) cannot suffice to support the strong coherence charge, without the aid of the virtually unsupportable premise that all possibilities have been exhausted. (That is, it is an elimination argument, that neither A nor B nor C nor D provides satisfactory ways of drawing the analytic/synthetic distinction.) Furthermore, Quine (1991) explicitly stops short of the strong coherence charge, stating there that the problem is not that it is incoherent to claim that (say) ‘All bachelors are unmarried’ is analytic, but rather that such trifles are of no philosophical interest or worth.

In any case, there is an argument for the strong coherence charge in Quine’s corpus, and it has been influential. (Cf., e.g., Boghossian & Peacocke (2000: 4): ‘Our own view is that Quine decisively refuted the idea that anything could be true purely in virtue of meaning.’) The core idea is that ‘no [statement] is true but reality makes it so’ (Quine, 1970: 10).¹⁶ Quine’s (1970: 10–12) argument in favor of this dictum runs as follows: The truth-conditions of any statement can be specified ‘as Tarski taught us’ via the schema: ‘S’ is T iff S. The disquoted S on the right hand side of the biconditional stands for a fact, an element of reality; and this is so whether the left-hand quoted ‘S’ expresses a contingency (such as ‘Quine speaks Portuguese’) or a triviality (such as ‘All Portuguese speakers speak Portuguese’). Therefore, no statement is true but reality makes it so.¹⁷ (Cf. Boghossian [1997: 334–37] for a very similar argument, inspired by Quine.)

While this argument has an unassailable air, I will argue that it commits a fallacy of equivocation. For ease of exposition, I will abbreviate Quine’s dictum (i.e., ‘no statement is true but that the facts make it so’) as ‘TVF’ (or ‘true in virtue of facts’), and the very idea that a statement might be true in virtue of meaning as ‘TVM’. The structure of the argument, then, is as follows:

P#1: TVF is unassailable.

P#2: TVF is incompatible with TVM.

Therefore, TVM is untenable.

There is a sense of ‘fact’ according to which TVF is unassailable, and a sense of ‘fact’ according to which TVF is incompatible with TVM. However, these are distinct senses of the term. There is no one reading that makes both premises true. Thus, this case against TVM is flawed by a fallacy of equivocation.

Consider first P#1, which Boghossian motivates with the following sorts of consideration:

After all, if a statement is known *a priori*, then it must be true. And if it is true then it must be factual, capable of being true or false. (1997: 334)

In this sense of ‘factual,’ to be factual is to be truth-evaluable—it is to make a claim, or to have truth-conditions. Note how heterogeneous the correlative conception of a fact must be, metaphysically speaking. For example, that no grandmothers are childless, that unicorns have one horn, that there is an even prime, that Germany is not in Asia, that hydrogen is less dense than gold, that justice is a virtue, that all humans are mortal, that there is no rhinoceros in this room, etc., are one and all facts, in this sense. The facts, in this sense, are as diverse in nature and status as are the indefinite range of things about which humans can think and talk.

This is the sense of ‘factual’ involved in Quine’s (1970: 10–12) Tarski-inspired argument described above. If the criterion for counting as factual is to issue in a truth when plugged into the schema: ‘S’ is T iff S, then it is not a very discriminating property. Since this notion of ‘fact’ does not correspond to any precise or homogeneous metaphysical category, I will call it the ‘Tarski-semantic notion of fact’.

Given this Tarski-semantic sense of fact, P#1 is uncontroversial, but P#2 is far from obvious. Consider, for example:

[1] No grandmother is childless.

It is, to say the least, not clear that there is any tension whatsoever between, on the one hand, the claim that [1] is T iff no grandmother is childless, and, on the other hand, the claim that one might be justified in believing [1] based solely on semantic intuition—that is, on understanding what it takes to be a grandmother—as opposed to based on empirical investigation or mystical Platonist intuition. That a sentence issues in a truth when plugged into the Tarskian schema (‘S’ is T iff S) is entirely irrelevant to questions about what it takes to understand or justify the proposition expressed. (Otherwise, that would spell trouble for a Tarskian approach to truth, one of whose key virtues is the clear distinction between the semantic concept of truth and, say, epistemic concepts like justification.) Thus, if we read ‘fact’ in this Tarski-semantic way, then, TVF is compatible with TVM. To say that a statement is

truth-evaluable does not rule out any possibilities as to what counts as understanding it, or how one might be justified in believing that it is true. On this Tarski-semantic reading of ‘fact,’ then, P#2 should be rejected.

Burge (2000: 16) affords another way to articulate this problem with P#2. When Leibniz or Frege contrast *a priori* truths of reason with *a posteriori* truths of fact, according to Burge, the point of the contrast is not that *a priori* truths are not factual, but rather that they are not *merely* factual. The claim is that such truths are not subject to refutation by contingent happenstance, not that they are entirely and categorically unrelated to contingent happenstance. So, [1], for instance, could be not *merely* factual, but not thereby non-natural, otherworldly. Grasp of its meaning might suffice to justify belief in its truth, but nonetheless the statement is still about our own flesh-and-blood grandmothers. So, again, it appears that P#2 is seriously flawed. A statement could be at once both factual, in the Tarski-semantic sense, while also being not *merely* factual, in the Leibniz-Frege sense. Thus, one and the same statement could be Tarski-semantic-factual and yet still be TVM.

There is a distinct, more metaphysically robust, conception of fact, given which P#2 fares much better. It is a long way from Tarskian schemas, though, closer to what Armstrong (1996) calls ‘truth-makers,’ or to Russell’s (1918) and Wittgenstein’s (1921) logical atoms. Facts in this sense are discrete mind-independent entities to which (typically: contingent, empirical) statements stand in some specific semantic relation (such as representation). Famously, Russell and Wittgenstein quarreled over whether it is a fact, in this sense, that there is no rhinoceros in the room. With respect to metaphysical worries about this variety of fact, negative existentials are just the tip of the iceberg. Russell (1918) seems to have never been able to convince himself of the existence of such general facts as that all humans are mortal, though he recognizes that his current views of meaningfulness seem to commit him to such entities. (Cf. Lewis [1998] for some related objections to Armstrong [1996].¹⁸) Russell (1918: Lecture 3) reports having ‘nearly produced a riot’ at Harvard in 1914 by arguing for the existence of negative facts. Surely, the claim that “‘There is no rhinoceros in this room’ is T iff there is no rhinoceros in this room’ would not have provoked such a reaction—even if the pragmatists would have disputed it.

In this metaphysical sense of ‘fact,’ there is definite tension between TVF and TVM.¹⁹ To be factual, in this metaphysical sense, is to represent a (typically: contingent, empirical) state of affairs; for any statement that represents a (contingent, empirical) state of affairs, there is reason to think that the ‘TVM’ label is probably inappropriate.²⁰ The price of this strategy of saving P#2 from imposing a false dilemma, though, is P#1. If we read ‘factual’ in this metaphysical sense, then TVF is eminently assailable, as is evidenced by Hume’s (1748) reasons for positing relations of ideas in the first place, by Russell’s (1918) struggles with negative and general facts, by Lewis’ (1998)

criticisms of Armstrong (1996), etc. In short, in order to make P#2 true, you have to think of TVF and TVM as mutually exclusive answers to an Armstrong-style demand for truth-makers. This issues an understanding of TVF that may well be interesting and even defensible, at least for a broad class of statements, but is certainly assailable. So, given this second, metaphysical sense of ‘fact,’ P#1 is controversial, not the sort of thing to which one can help oneself on the strength of vague slogans, or disquotational schemas.

To sum up: Quine (1970: 10–12) and Boghossian (1997: 334–37) provide clear illustrations of a prevalent line of thought that is widely but mistakenly thought to spell the end for TVM. The semblance of a compelling case against TVM depends on taking ‘factual’ in the Tarski-semantic sense in P#1 while taking it in the metaphysical-truth-maker sense in P#2. Once we recognize and guard against this slide, either P#1 is extremely contentious (if we adhere to the metaphysical-truth-maker sense of ‘factual’) or else P#2 is false (if we adhere to the Tarski-semantic sense of ‘factual’). So, there is no cogent case against TVM forthcoming down this avenue.

Given that these arguments against the coherence of TVM are lame, and provided that we have an account of analyticity that unequivocally rejects the notion that meaning has supernatural truth-making powers, analyticity may yet hold some promise to provide a compelling account of justification, for at least some of our universal, necessary knowledge.

[§]

One important and pertinent point which Quine (1951) is commonly taken to have established is that ‘analyticity’ admits of no satisfactory non-circular, reductive conceptual analysis. That is, analyticity cannot be reduced to, or defined purely in terms of, some other less obscure notion. Hence Quine is often taken to have demonstrated that there is something deeply suspect about appeal to this notion. It might even be taken to be another sort of coherence argument: analyticity can play no role in a rigorous philosophical theory, because there is no acceptable, non-circular definition of the notion.

For example, if we had a satisfactory criterion for synonymy, then we could get from there to a satisfactory definition of analyticity. (That is, if it were possible to conclusively establish that ‘bachelor’ is synonymous with ‘unmarried man,’ then it could be demonstrated that it is analytic that all bachelors are unmarried men.) Vice versa, if analyticity were unproblematic, then a crisp, clear definition of synonymy would be forthcoming. But one cannot help oneself to synonymy in definiendum analyticity, so the objection goes, since the definiendum is equally as obscure and suspect as the definiens.

First: Note the role that meaning-skepticism plays here. This line of argument draws force from the sentiment that meaning is indeterminate, inscrutable, and hence that synonymy is deeply suspect. Again, as pointed out

above, this line of argument is considerably less weighty in the wake of the cognitive revolution.

Second: Of course, lots of (probably, most?) concepts cannot be reduced to others without remainder.²¹ Why think that they ought to be? Why hold ‘meaning’ to a higher standard than the rest of the lexicon? This elementary status should not be all that surprising, and is hardly a defect. In general, there is no good reason to hold that that renders them suspect, or second-rate. For analyticity, as for any other concept, showing that it cannot be reduced to any other concepts is far from sufficient for showing it to be incoherent (without the aid of a premise, or presumption, that there is something deeply suspect or second-rate about meaning).

[§]

On these grounds, I take coherence of analyticity to be a reasonable tenet. Assuming that we all concede the weak coherence charge, the live coherence dispute is then over the question of whether the analytic/synthetic distinction is worth further investigation. There is weighty philosophical-historical precedent, as well as the brute data, in favor of further pursuit of a comprehensive, satisfactory conception of analyticity. There are no compelling arguments to the contrary of which I know—provided that we are careful about boundaries between metaphysics, semantics, and epistemology, and pay clear-headed heed to fine semantic distinctions between questions that pertain to the conventional links between S and p and questions that pertain to the truth-conditions of p.

Some prevalent arguments against analyticity do not in fact bear upon its coherence or worth. The notion has not been proven incoherent, and it still may well be our best hope for an adequate, non-obscure solution to the problem of *a priori* knowledge—among other things.

A reason to retain analyticity is that the brute data is not going away. Given the lack of other viable accounts of universal, necessary knowledge, this avenue is still eminently worth further exploration.

There remains of course the massive and daunting challenge of revisability. This is a good segue to chapter 4 as that issue is also deeply pertinent to the notion of *a priori* knowledge. On, then, to chart the *a priori*.

NOTES

1. Cf. Cartwright (1983) for thorough recent development of this line of thought.
2. Cf. especially Kripke (1972: Lecture II). Roca-Royes (2011) is a good general discussion of these issues.
3. Cf. Lewis (1986) for development and defense of an anti-haecceitist approach.

4. Cf. Kripke (1972, 41–53; 1980, 15–20).

5. See Boghossian (1997) and Sullivan (2008) for further justification for taking the statement as the primary bearer of analyticity.

6. Williamson (2007, 2014) and Hawthorne (2012) have recently argued for some such conclusion about *a priority*; but even still, there remains a vast gap between it and analyticity (let alone necessity) in this respect.

Note that I do not mean to imply that all or even most philosophers dismiss analyticity as a worthless notion. To the contrary, consider for example Grice's (1987: 344) claim that the analytic/synthetic distinction is 'one of the most important topics in philosophy, required in determining, not merely answers to particular philosophical questions, but the nature of philosophy itself'.

7. I argue this at length in Sullivan (2008). Grice & Strawson (1956) made an early, forceful case in favor of this claim, and for a varied sample of subsequent supporters of this claim cf. Fine (1994), Katz (1997), and Gertler (2002).

8. For some recent statements cf. Sullivan (2003a) and Cappelen & Lepore (2005). Here is a classic statement of the core idea from Frege (1892b: 46):

Nowadays people seem inclined to exaggerate the scope of the statement that different linguistic expressions are never completely equivalent, that a word can never be exactly translated into another language. One might perhaps go even further, and say that the same word is never taken in quite the same way even by men who share a language. I will not enquire as to the measure of truth in these statements; I would only emphasize that nevertheless different expressions quite often have something in common, which I call the [meaning], or in the special case of sentences, the [proposition]. In other words, we must not fail to recognize that the same [meaning], the same [proposition], may be variously expressed. ... It is possible for one sentence to give no more and no less information than another; and, for all the multiplicity of languages, mankind has a common stock of [propositions]. If all transformations of the expression were forbidden on the plea that this would alter the content as well, logic would simply be crippled; for the task of logic can hardly be performed without trying to recognize the [proposition] in its manifold guises. Moreover, all definitions would then have to be rejected as false.

As with some other excerpts, I have freely substituted some of Frege's terms (e.g., 'sense', 'thought') with similar terms (e.g., 'meaning', 'proposition'), in order to better fit this excerpt with the broader discussion.

9. Cf. Sober & Hylton (2000) for development of a similar theme. For an introduction to the cognitive revolution, cf. Pinker (2011).

10. For discussion see Munsat, ed. (1971: 19–20).

11. Again, as with some other excerpts, I have freely substituted some of Coffa's terms (e.g., 'concept', 'representation') with similar terms (e.g., 'meaning', 'proposition'), in order to better fit this excerpt with the broader discussion.

12. See Sullivan, ed. (2003) for extensive treatment of this allegation. The core idea is that, whereas traditional logic viewed the content of any significant proposition along the lines of:

[All/some/no] S [is/is not] P

Various nineteenth-century (and hence post-Kantian) logicians discovered that this does not capture the correct logic of various sorts of propositions.

13. See especially Frege (1884) and Quine (1951) for seminal statements, and Coffa (1991: Part 1) for discussion.

14. For critical discussion of Russell (2008) cf., for example, Wilkfors (2008), Boghossian (2010).

15. Quine is commonly thought to be a leading proponent of the stronger charge against the coherence of metaphysical necessity. However, Quine's target is not the intelligibility of metaphysical necessity per se, but rather the prevalent early-twentieth-century empiricist idea that all necessity reduces to analyticity, and that that affords a decisive solution to Plato's problem. For good discussions cf. Kaplan (1986), Neale (1990: Ch.4), Marcus (1991).

16. Here and throughout, to fit the terms of the present discussion, I use 'statement' in place of Quine's nominalist adherence to the term 'sentence'.

17. In additions to Boghossian (1997), Cassam (2000) is another recent writer who is clearly influenced by this argument. The case of Boghossian is complicated, though. He distinguishes between metaphysical and epistemological conceptions of analyticity, and argues that while Quine does present a conclusive case against the metaphysical conception, it does not affect the worthwhile and valuable epistemological conception. While I agree with much of what Boghossian has to say about what he calls 'the epistemological conception of analyticity', there is much that I disagree with here. First of all, analyticity is a semantic notion, and so should be firmly distinguished from both metaphysical and epistemological concepts. Second, as mentioned above (in §2.1), I think that what Boghossian calls 'the metaphysical conception of analyticity' is completely a straw target anyway—no one that I can think of believes that our meaning p by S makes it the case that p .

18. Note that something like Quine's (1970: 10) dictum that no statement is true but reality makes it so explicitly plays a role in Armstrong's (1996) case for positing truth-makers.

19. It is in this second, metaphysical sense of 'fact' that the approach to *a priori* knowledge associated with Hume (1748) and Ayer (1936) is not unfairly glossed as the view that *a priori* knowledge is devoid of factual content. (I take it that Hume and Ayer are saying something different from the above-discussed view that *a priori* knowledge is not *merely* factual.) That is, what Hume and Ayer claim is that the likes of:

[2] All squares have four sides.

do not rule out any contingent empirical possibilities, and this explains why they are not subject to refutation by contingent happenstance. Obviously, but nonetheless crucial for present concerns, the Hume-Ayer claim is certainly not in the slightest tension with anything along the lines of: [2] is T iff all squares have four sides. Hume and Ayer have no reason or inclination to deny that such statements are Tarski-semantic-factual.

20. Though this does accord with Hume (1748), and the wealth of tradition, it might run afoul of Kripke (1972) on the contingent *a priori*. In any case, as I explain immediately below, my defense of TVM in no way depends on siding with Hume against Kripke on the contingent *a priori*, because P#1 is false on this second understanding of 'fact' anyway. (There is more on the contingent *a priori* in Part IV.)

21. Compare the remarks about 'truth' (or 'good', or 'art') in §2.1.

Chapter 4

A Priori Justification

§4.1: FLESHING OUT THE CONCEPT '*A PRIORI*'

At first pass, *a priori* knowledge is knowledge whose justification does not depend on sensory evidence. In the classic sense of the term 'empirical,' that is, 'via the senses,' the *a priori*/*a posteriori* distinction corresponds to the non-empirical/empirical distinction.

This particular usage of the terms may just date back to Kant or Leibniz, but the distinction they mark is as old as philosophy itself: the existence and nature of *a priori* knowledge is an absolutely central, core issue throughout the history of philosophy. For one thing, as the cases of at least Plato, Descartes, Kant, Frege, and Quine illustrate, some seminal, original philosophical systems and stances have grown out of work on the nature of the *a priori*. For another thing, a philosopher's view about the existence and nature of the *a priori* is inextricably linked to not only various other issues in metaphysics, semantics, and epistemology, but also to their conception of the discipline of philosophy as a whole—including in particular its proper methodology and its scope or range (cf. Peacocke [2006]).

I began §1.1 above by motivating the claim that 'one cannot steal one's own property' is a strong candidate example of being justified *a priori*. For another plausible candidate, compare what it would take to be justified in believing the following:

1. Squares have four sides.
2. Neptune has four moons.

For both [1] and [2], understanding the sentence affords a grasp of *what would have to be the case* for it to express a truth. However, for the case

of [1], this understanding also and thereby *justifies* the belief that what is expressed is true. One need not to take a poll to find out how many people agree, or devise a variety of experiments to test whether one could succeed in constructing a three- or five-sided square, in order to be justified in believing that [1] is true. Although one is justified in believing [1], this justification has nothing to do with that kind of empirical evidence. This might be taken to be a basic, straightforward example of immunity to counterexample, grounded in (some or other variety of) non-empirical justification.

Not so for [2], in which case understanding it does not come remotely close to providing justification for believing that what it expresses is true. Even though I know exactly what [2] means, I have no justification as to whether or not it is true. (Note that it is far from clear that a correlative claim could coherently be made about [1].) Case [2], it seems, could only be justified empirically—akin to cases like ‘There is an apple pie baking nearby’ or ‘This piano needs tuning,’ discussed in §1.1 above.

Further examples of beliefs that are justified *a posteriori* include my beliefs that it is not currently raining here now, that it rains more frequently in England than in Arizona, and that it is hard to find a direct flight from England to Arizona. Reason, by itself, is not able to afford a grasp of the truth of such statements. (One needs to stick one’s head out the door, or to travel a bit, or to ask meteorologists or travel agents.) Some of the things that some philosophers have argued are justified *a priori* include elementary truths of logic and mathematics (e.g., ‘No number is both even and odd’), and certain fundamental truths about human beings (e.g., Plato holds that it is knowable *a priori* that humans have immortal souls; he and others have argued that many moral truths [such as that one ought to keep one’s promises] are also knowable *a priori*). There are good reasons to doubt that empirical justification could suffice to support these claims; so, it seems, one must either be skeptical that we know them, or else hold that there is *a priori* justification.

More generally, most philosophers have held that not all of our knowledge can be seen as empirically justified, and so posit the category of *a priori* knowledge. In particular, from at least Plato on down, the prevalent view is that some of our knowledge is simply immune to counterexample, and that such knowledge could not possibly be justified empirically.

As discussed in §1.2, philosophers who are comfortable with talk of *a priori* knowledge tend to talk also of rational intuition—that is, the non-empirical faculty of mind involved in *a priori* justification. Opponents counter that this appeal to rational intuition is merely a label for the problem, not a solution to it, that such appeals to mysterious inexplicable faculties are too obscure to be of any help in rigorous epistemology. The rationalist then counters that no epistemology which denies that there is *a priori* knowledge is at all adequate to the task of a satisfactory account of our actual human

knowledge. Variations on this theme of the obscurity objection to rationalism versus the adequacy objection to empiricism make up a main thread running through the history of Western philosophy.

Some influential variations within this debate are focused on the notion of analytic truth and semantic intuition, and their potential to provide a non-obscure but adequate solution to Plato's problem. Relatedly, many significant recent episodes have involved the development of understanding-based accounts of *a priority*, as distinct from acquaintance-based accounts. These themes will be revisited and further developed, especially throughout Part III below.

[§]

Even more so than the cases of 'necessary' and 'analytic,' there are several different—overlapping but nonetheless prima facie non-equivalent—objects to which the term '*a priori*' is attributed. To what, precisely, does the term '*a priori*' appropriately apply? It is not uncommon to find the label '*a priori*' affixed to, among other things: (i) concept, (ii) inference, (iii) proposition, (iv) truth, (v) knowledge, and (vi) justification. Some useful refinements will emerge from discussing some of the relative merits of these notions.

I will take the notion of an (i) *a priori* concept first. An important refinement prompted by consideration of this notion is that *a priority* is distinct from innateness—despite there being some broad and deep connections between these two notions. That is: the claim that one's knowledge that P is justified *a priori* does not entail that no experience is required to acquire the concepts needed to grasp P. It is not the claim that all of the concepts involved are innate, or in every sense self-evident, that they can be acquired without causal and communicative interaction with one's environment and linguistic community. (Indeed, I am suggesting that 'One cannot steal one's own property' is justified *a priori*, and of course a range of experiences are required in order to acquire the concept of property.) One may need experience to get the concepts, which one can then exploit in cultivating *a priori* knowledge. What matters for the *a priori/a posteriori* distinction is whether interaction with, experience of, the specific object of the specific belief is sufficient to justify belief in its truth. Even if I may need to undergo some kinds of experience in order to acquire the concept of number, I do not need acquaintance with every number, to consider each case individually, in order to know that all numbers are either even or odd. In contrast, I do need some kind of contact with the weather in England and the weather in Arizona, by personal experience or the testimony or expertise of others, in order to know that it rains more in England than in Arizona.

Kant (1781) and Frege (1884), among others, are very clear on this point. *A priority* is not innateness; claims of *a priority* are compatible with various

sorts of important role for experience, when it comes to conditions for acquiring various sorts of concept. What matters is the justification for the specific item of knowledge in question, not its psychological cause. ‘*A priori*’ is a term of normative epistemology, which applies to relations between concepts; it is not a psychological term, which applies to how in fact this particular person came to have the particular conceptions, or beliefs about those concepts, which they have.

To be sure, the Modern rationalists who were such influential friends of *a priori* were also quite up to their necks in claims about the innate, God-giveness of much of human knowledge. Plato is another influential example of someone who believes strongly in both rational intuition and innateness. (Famously, Plato taught that all knowledge is recollection—cf., for example, the *Meno*.) However, the fact that lots of smart people believed in both A and B does not entail that $A = B$. There are important differences between *a priori* and innateness, and that is one of the main reasons why I will not speak at all in terms of (i) above, the notion of an *a priori* concept.

Now as for (ii), the notion of an *a priori* inference: to be sure, this is an important concept historically and conceptually, tied up as it is with such notions as logical truth, entailment, validity, etc. However, this present work does not directly venture into the philosophy of logic, in any sustained way. First and foremost, inferences pertain to relations between propositions; whereas my primary interest here is in properties of propositions. (Generally, relations between concepts are also and thereby properties of propositions; but questions about relations between propositions arise at another level of inquiry.) In any case, *a priori* will not be directly applied to inferences below either.

So, what then of (iii) *a priori* proposition or (iv) *a priori* truth? ‘*A priori* truth’ is certainly a common locution in philosophy, and this notion is closely related to the core epistemological distinction that is our present focus—for if a piece of knowledge does not depend on experiential evidence, then it seems that there must be something special and distinctive about the proposition which is the object of knowledge. So, an *a priori* truth would be something which can be known to be true without sensory evidence.

However, the question of the distinctive status of such truths or propositions seems rather clearly to be, at least to a considerable extent, a semantic question; whereas here we are stalking a distinction that is epistemic (at least: as purely so as is possible). Henceforth, I will avoid the notion of ‘*a priori* truth’—at worst, it is a potentially misleading label for analytic truth, at best it is a shorthand for one of the more purely epistemic notions to be developed below.

Likewise for (iv), ‘*a priori* proposition.’ While I am stalking a property of propositions, it is an epistemic property, and not a semantic one. So, the notion of an *a priori* proposition is not the optimal label for the target. ‘*A priori-ity*’ does not apply, first and foremost, to semantic notions like ‘meaning’ or

‘proposition,’ but rather to epistemic notions like ‘belief,’ ‘justification,’ ‘warrant,’ ‘knowledge.’

The also common notion of (v) ‘*a priori* knowledge’ is an improvement, in this respect. It is more clearly not semantic and quite distinct from analytic truth. The idea here is that there are (at least) two different varieties of knowledge—knowledge that depends on experiential evidence, and knowledge that does not. Again, this epistemological distinction is clearly evident in Plato’s work, and throughout much subsequent work in epistemology. Many have held that an adequate epistemology must include both, on the one hand, the (non-empirical) knowledge we can arrive at solely via rational reflection, solely via the faculty of pure reason, and, on the other hand, the (empirical) knowledge that we attain via causal interaction, through sensory channels, with specific individuals and states of affairs in our environment.

Still and all, ‘*a priori* knowledge’ is at best a convenient shorthand for the more cumbersome but correct ‘belief which is justified *a priori*.’ For one thing, I wish to ward off any suggestion that ‘*a priori* knowledge’ essentially involves rational intuition, in any substantive sense, because I certainly want to leave room on the playing field for understanding-based accounts—that is, those who are dismissive of rational intuition (as obscure and unhelpful) but not thereby dismissive of the view that no adequate epistemology can avoid the notion of *a priori* justification. So I want to be clear and upfront that my use of ‘*a priori*’ does not essentially connote a distinct and possibly supernatural faculty of mind. For example, at least arguably, semantic intuition might provide the grounds for *a priori* without appeal to, or need of, any such distinct faculty of mind.

Second, when it comes to one of our ongoing themes of the distinction between framework- (and vehicle-) indifferent metaphysical questions on the one hand and framework-relative semantic and epistemic issues on the other hand, there is a big difference between ‘*a priori* knowledge’ and ‘*a priori* justification.’ It rests on the consideration that ‘knowledge’ is a factive term—that is, ‘A knows that P’ entails that P is true. Hence, staking a claim to knowledge involves progressing beyond the framework-relative bounds of epistemology; it makes a claim about the world, in addition to a claim about an agent. So, ‘X is justified *a priori*’ is a weaker, more purely epistemic claim, as compared with the bolder ‘X is known *a priori*.’ Just as something could be empirically well-justified but yet not true (e.g., the sun revolves around the earth, all swans are white), it is at least a coherent possibility for something to be justified *a priori* but not true. As these boundaries are drawn and developed in this present work, *a priori* justification is the focal epistemological issue.¹

The fundamental distinction is the following one between (vi) *a priori* and *a posteriori* justification. One core central issue throughout the history of epistemology has concerned exactly how to distinguish between mere belief

and genuine knowledge. It is more or less universally conceded that the concepts of truth and justification must be involved in a satisfactory account of this fundamental distinction. That is, for a belief to count as knowledge, at the very least it must be both true (i.e., 'X knows that P' can only be true if P is true) and justified (because, in general, we do not count people who are right completely by accident as possessing knowledge). There are lots of involved disputes over exact details, and further conditions.²

So, assuming that knowledge is at least justified true belief (and leaving open the question of what else might be required of an adequate definition of knowledge), it has seemed to Plato and most philosophers since that we need to posit at least two different kinds or types of justification. Henceforth, the term '*a priori*' will always be understood as applying to (vi) justification.

[§]

The *a priori*/*a posteriori* distinction is fundamentally a distinction between two characteristic ways in which beliefs are justified. If the justification is that a belief must involve appeal to sensory evidence, then it is *a posteriori*. Typically, *a posteriori* justification rests on a causal story about relations between an agent and the object of knowledge; in contrast, *a priori* justification has always resisted smooth incorporation into any such straightforward, scientifically tractable model. This is one of the reasons why empiricists and skeptics have always been wary of appeals to *a priority*.

Why should we believe in *a priori* knowledge? Next, I will give an overview of some related, classic arguments, to give a sense of the reasons in favor of *a priority*. They all date back at least to Plato. The arguments are interconnected, and mutually supporting. They are not always distinguished, and may well be seen as three different aspects of the same underlying phenomenon (i.e., Plato's problem). The core idea is a sort of *poverty of experience*: all of our actual knowledge cannot possibly be *a posteriori* inductions from experience, because our actual knowledge outruns, in principled ways, our experiences. Hence, while much or most of our knowledge can be understood as justified by causal connections to the ambient environment (including testimony as a special, mediated case), not all of our knowledge can be seen as so justified. Because our knowledge outstrips our experience in this way, our faculty of reason must be seen as contributing to, supplementing, or structuring what we learn from experience.³

I will refer to three related strands of this line of argument in favor of *a priority* as: (i) experience is particular but some knowledge is general, (ii) empirical justification can only afford inductive generalizations, and (iii) some knowledge just simply glows with luminous certainty.

The first (i.e., experience is particular but some knowledge is general) strand of the argument is clearly evident in Plato's dialogues. The world

experienced by the senses is a world of Heraclitean flux, constantly changing from moment to moment. Still we, in fact, manage to know a lot of stuff that is indifferent to empirical and psychological contingencies. Therefore, we must have this distinctive and remarkable faculty of mind—that is, rational intuition. Whether we are talking about logic and mathematics, moral truths, or other domains, this strand of the argument has played a key role in the history of epistemology.

This core strand within the poverty-of-experience line of argument is focused on the remarkable differences between the contents which make up our experiences and the contents which compose some of our (presumed) actual knowledge. And it is taken by many (including Plato) to support not only rational intuition but also metaphysical realism about the objects of our *a priori* knowledge. Many mathematicians, for example, find these considerations compelling, and to warrant not only epistemological but also metaphysical conclusions. Consider, for example, the following sentiments expressed by Russell (1912: 100), from the ranks of mathematicians who defended the indispensability of rational intuition, describing the otherworldly objects of our *a priori* knowledge:

The world of universals ... is unchangeable, rigid, exact, delightful to the mathematician, the logician, the builder of metaphysical systems, and all who love perfection more than life. The world of existence is fleeting, vague, without sharp boundaries, without any clear plan or arrangement.⁴

The second strand of poverty-of-experience argument for *a priori* may well just be the distinctive way in which the first strand of the argument gets formulated in response to a certain kind of empiricist agenda. Classic instances include Leibniz' (1704) response to Locke (1690), Kant's (1781) response to Hume (1748), and Frege's (1884) to Mill (1831). You only get inductive generality, not necessity, from experience. Given that we do attain knowledge of (at least some) necessities, it follows that at least some of our knowledge is not merely a matter of what the mind has soaked up or observed via the senses, but must rather be a non-empirical matter—that is, grounded in rational intuition. *A posteriori* faculties will ground the claim that all events that I have yet encountered have been caused, but not that it is necessary that all events have been caused; or the claim that every murder I have ever encountered is morally unjustified, but not the fully general claim that murder is unjust. Once you stir in the claim that we do in fact have knowledge of these general necessities, this is a seminal influential argument for the rationalist position.

A textbook example for making this point is the well-confirmed but ill-fated inductive generalization that all swans are white. This was taken to be

a completely exceptionless law by Europeans, until, surprisingly, black swans were discovered in Australia. They were wrong about swans, and in retrospect the judgment should have been qualified to: all the swans which we have ever observed have been white, but there is nothing contradictory or impossible about a non-white swan. Now, here is the rationalist kicker—this situation could not possibly occur for ‘all squares are four-sided,’ or for ‘ $2 + 2 = 4$.’ Explorers of the nether reaches are not going to discover three- or five-sided squares; or discover a place in which $2 + 2$ yields anything but 4. Hence, there is a difference between defeasible inductive generalizations like ‘all swans are white’ and statements which are immune to counterexample; and—the argument continues—only by positing rational intuition (in one form or another) can this evident difference be explained.

As for the last strand of poverty-of-experience line of argument in favor of *a priori*, ‘luminous certainty’ is a term which occurs in Leibniz (1704) to denote a special phenomenal quality appropriate to certain self-evident unassailable contents. Some knowledge is just bulletproof, solidly and firmly immune to counterexample. Such knowledge just glows with a luminous certainty. To understand some privileged claims is to be sure that they are true. And, so the argument continues, there is nothing within the *a posteriori* evidence-gathering faculty which could ground or underwrite that. Only a non-empirical source of justification will do.

All three interrelated strands of argument are taken by some (from Plato on down) to support not only rational intuition, but also metaphysical realism about the objects of our *a priori* knowledge. Many are tempted to continue from metaphysical realism on into an anti-scientific mysticism (i.e., here we have an obvious and undeniable example of something which lies beyond the bounds of scientific explanation). However, positing rational intuition does not entail any such metaphysical realist or mystical views. These are considerably strong arguments in favor of the positing of a non-empirical source of justification (provided of course that one grants the anti-skeptical premise that we do manage to obtain knowledge of universal, general truths); but per se they are much weaker arguments for any specific metaphysical position. Many of the moderate varieties of both rationalism and empiricism are precisely attempts to concede the force of these epistemological arguments without drawing any unwarranted metaphysical or mystical conclusions.

[§]

To sum up then: ‘*A priori*’ applies first and foremost to a putative kind of, or source of, justification. Humans have *a priori* knowledge only if it is not the case that all of our knowledge is justified empirically. Rationalists are philosophers who are enthusiastic about *a priori* justification, and ‘rational intuition’ is a common but vague term for this non-empirical source of justification.

Empiricists are opposed to any substantive form of rational intuition, and the recent naturalist (or radical empiricist) strand within empiricism seems to be opposed to any and all forms of *a priori*.

As opposed to the metaphysical notion of necessity, or the semantic notion of analytic truth, *a priori* is an epistemological concept. The bone of contention between proponents and opponents of *a priori* concerns what is required for an adequate account of human knowledge. (Will *a posteriori* justification suffice, or not?)

This is a central distinction within the history of philosophy, and much interest has lain in trying to explain these two different kinds of knowledge. What is this psychological and empirical indifference which seems to be a mark of some of our knowledge? What explains the evident fact that some of our knowledge is immune to counterexample?

§4.2: REFINEMENTS AND ELABORATIONS

While the notion of *a priori* justification is more or less as old as philosophical reflection on human knowledge, philosophical accounts of *a priori* are a much more recent phenomenon. Historically, instances of employing the concept of *a priori* are common, but attempts to define the notion are rare. Plato gives some suggestive but vague descriptions of the phenomena, and many of his followers were inclined to view *a priori* knowledge as mystical, non-discursive, not amenable to analysis or reductive definition. In the work of the Modern rationalists (such as Descartes, Spinoza, Leibniz), one finds *a priori* integrally bound up in the nexus of a certain distinctive kind of Enlightenment-scientific-theistic world-order (i.e., *a priori* knowledge is a gift from God, the benevolent omnipotent watchmaker, to allow us a glimpse of the divine geometrical plan). Needless to say, many philosophers today are skeptical about many different aspects of that world-view.

Burge (2000: 13) alleges that Leibniz was the first to attempt to give a rigorous definition of the notion of *a priori*; and many current discussions of the notion take Kant's work (on which Leibniz was a principal influence) as its starting point. I will follow in this vein, starting this portion of our investigation from Kant's Leibnizian definition of the *a priori*, and investigating a couple of questions which are raised thereby.

Kant's (1781) most common, official definition is negative: *a priori* knowledge is knowledge that is *not* dependent on experience. This immediately gives rise to two related questions: First, what exactly does 'experience' mean here? Second, is it possible to give a purely positive definition of *a priori* knowledge? I will explore both of these in a preliminary way here. They are so central that the discussion of them will continue on into Part IV.

So, what exactly does ‘experience’ mean, when we classify *a priori* knowledge as independent of experience? Presumably not just ‘current, operative sensory experience,’ as that would count things justified via memory (e.g., ‘it rained yesterday’) and introspection (e.g., ‘I am definitely not repressing my frustrations’) as *a priori*; and yet there is reason to classify at least some of such things on the *a posteriori*, empirical side of the divide. (They can be based on causal transactions, and can be deeply fallible, and so lack the distinctive sort of privilege that we have with the paradigm sorts of cases which we want to classify as *a priori*.) However, we do not want to go too far in the other direction, and characterize experience so broadly as to include any kind of subjective psychological process; for the most extreme Platonic ‘the very light of the mind in her own clearness’ sorts of purely mystical cases (discussed in §1.2) would count as experience in this sense. And if the insights gleaned from a mystical Socratic trance⁵ do not count as *a priori*, then nothing does!

Questions regarding whether introspection counts as experience (or, relatedly, whether self-knowledge counts as *a priori* or empirical) raise vexing complications. Is ‘I am in pain’ *a priori* or empirical? What about ‘I am currently conscious’? Or how about Descartes’ famous unassailable *cogito*, ‘I think, therefore I am’? We certainly enjoy a kind of privileged access to our own subjective mental states, but at the same time there is a degree of fallibility, a possibility of error or of self-deception, which marks off at least some such cases as distinct from our paradigmatic *a priori* cases. We might put it this way: I have privileged access to what my mental states are, but not to whether they are knowledge—as to what their content is, but not as to whether that content is true.⁶

In any case, what is needed, it seems, for a notion of *a priority* which is sharp, defensible, interesting, and not out of touch with its broad historical roots, is a sense of ‘experience’ which counts some but not all introspective data, subjective psychological processes, as experiential. This is not an easy thing to define. For example, one fairly compelling way to draw the key *a priori*/empirical difference would be to do so in terms of the presence, or absence, of a causal link between knower and known. The idea would go something like: A’s knowledge that P is *a posteriori* iff A’s justification for thinking that P depends upon a causal link between A and P; A’s knowledge that P is *a priori* iff there is no such causal justificatory link between knower and known.

There certainly is something to this way of drawing the contrast—particularly if we take Plato’s combination of acquaintance-based, rational intuition and metaphysical realism as the typical proponent of *a priority* (because there are no causal connections between human brains and Platonic universals). It seems to fit well with the paradigm cases of *a priori* and empirical knowledge. However, on reflection, surely this would put way too much into the *a posteriori*, empirical basket. For starters, all self-knowledge cases (e.g., I exist, I am conscious, I

am hungry) seem to involve causal relations to brain events. Indeed, isn't every mental event, including a mystical Socratic trance, causally related to brain events? Is everything, then, *a posteriori*, excepting on a scarcely tenable extreme dualism which denies any causal relations between minds and brains? So, again, what initially seemed to be a straightforward and promising way to distinguish *a priori* from empirical turns out to be rife with complications.

This challenge as to the exact sense of 'experience' which is involved in the notion that *a priori* knowledge is independent of experience will be further discussed in Part IV (at §7.4). One thing which we can say with a fair degree of security is that it is not clear that any clean, crisp way to define the *a priori/a posteriori* divide will suffice to capture all historically significant work on the *a priori*.⁷

[§]

What, then, of the second question, as to whether or not *a priori* justification could be defined in a positive way (as opposed to merely as the absence of something, such as empirical experience, or a causal connection between knower and known)? Well, the notion of rational intuition might seem to provide a candidate, the idea being that A's knowledge that P is *a posteriori* iff it is justified empirically, and *a priori* iff it is justified via rational intuition. Fair enough, as far as it goes; the problem here being that it does not go very far at all. As empiricists have long complained, rational intuition thus understood is merely a label for the problem, not a solution to it. The burden of coming up with a satisfactory explanation of *a priori* justification is merely recast as the burden of coming up with a satisfactory explanation of rational intuition.

The appeal to semantic intuition, or the move from acquaintance-based to understanding-based accounts, is seen by many as a step forward at this point, as it seems to have some promise to flesh in what is required from the above vague appeal to rational intuition, without taking on the obscure epistemological or metaphysical commitments associated with some extreme varieties of rationalism. The task, then, is to flesh in some such strategy, of grounding *a priori* justification in semantic intuition or understanding, in a satisfactory way. Here the price of this way of addressing the obscurity objection is that the adequacy objection looms larger. The question thus prompted is whether semantic intuition can get beyond mere trifles and trivialities, to do any of the heavy lifting which a solution to Plato's problem requires. It is a main task of Part III below to show how the notion of the constitutive *a priori* is a great leap forward, when it comes to these problems.

[§]

As with the cases of 'necessary' and especially 'analytic' (cf. §3.2) above, one main aim of this chapter is to disavow some of the traditional

associations with the term ‘*a priori*.’ While it is important that the concepts employed herein are recognizably in touch with their historical roots, it is neither desirable nor indeed possible to retain everything that anyone has ever claimed about the *a priori*, if I am to defend specific and focused theses about the notion. (This present task of pruning some historical associations with ‘*a priori*’ will overlap somewhat with the task undertaken in the previous section of determining to what to which the term ‘*a priori*’ first and foremost applies. At a few points below where that is so, this present discussion can be quicker and neater because of that previous work.)

Traditionally, ‘*a priori*’ has been tightly linked to such notions as (i) innateness, (ii) generality, (iii) infallibility, (iv) privileged access, and (v) unrevisability. These links have come under much scrutiny and pressure, in recent decades. To some extent, these pressures also affect necessity and analyticity; but they are most pertinent here. If immunity to counterexample is assailable anywhere, it is here on the epistemic front (as we are obviously limited, fallible epistemic agents). Furthermore, the idea of self-consciously weaker, less ambitious varieties of *a priori* has been more explicitly worked out in the literature here (as distinct from the cases of analytic or necessary truth).⁸

(i) The putative links between *a priori* and innateness have already been addressed above in §4.1, in giving reasons against employing the notion of an *a priori* concept. While lots of influential philosophers (including in particular both Plato and the Modern rationalists) tend to endorse both innateness and *a priori*, while not distinguishing between these two different sorts of independence from experience, Kant and Frege, among others, are clear about the importance of this distinction. A judgment might be justified *a priori*, even though there is nothing innate about any of the contents which compose its content.

(ii) As for the relations between *a priori* and generality, this complicated question has also been broached. Burge (2000) discusses this question, in the course of which he distinguishes two broadly overlapping notions of *a priori* running through modern epistemology: the first, championed by Leibniz, takes generality to be the key criterion for *a priori*, and the second, associated with Kant, takes independence from experience to be the crucial mark. These two conceptions would overlap, over a broad range of paradigm cases; but insofar as there could be *a priori* knowledge of specific phenomena (e.g., Descartes’ *cogito*), then the two conceptions would not be extensionally equivalent.⁹

One important present point is exactly one which came up above, in our discussion of the relevant sense of ‘experience’ in the definition of *a priori*: namely, if we want to count any cases of self-knowledge or of introspection

as *a priori*, then generality is not an exceptionless mark of *a priori*. (There is more discussion of this point in Part IV.)

(iii) Infallibility was certainly one of the traditional marks of the concept of *a priori*, and indeed this aspect of immunity to counterexample has been, historically and conceptually, one of the main reasons why the notion is of such enduring philosophical interest. This is of a piece with the ‘luminous certainty’ aspect of *a priori*. Some knowledge (about mathematics or morality, say) just seems to be completely unassailable.

The view to be developed herein will instance a complex stance on this front. *A priori* justification, similar to analyticity, will turn out to be framework-relative and hence revisable. However, there is still a sense in which immunity to counterexample is retained. Any case in which what was previously taken to be *a priori* subsequently gets rejected involves conceptual evolution, and consequently a change of framework. (For example, Democritus’ belief that atoms are indivisible was not contradicted by what we in the twentieth century call ‘the splitting of the atom.’ Rather, the term ‘atom’ has undergone conceptual evolution. Cf. §6.3.)

As distinct from (though consistent with) this framework-relative revisionism about the semantic and epistemic modalities, some recent philosophers have favored a fallibilist stance toward the *a priori* (cf., e.g., BonJour [1998], Casullo [2003]; and see note 1 above on the putative ‘defeasibility requirement’). Fallibilists about *a priori* hold that, just as one can be justified *a posteriori* in believing something that happens to be false (e.g., the sun revolves around the earth), so too one could have a mistaken false belief that is nonetheless (more or less well) justified *a priori*.

The constitutive *a priori* view is compatible with this brand of fallibilism; in many respects to be developed, it can be seen as bolstering and extending this approach. Of course, fallibilism too undermines the kind of easy transition between *a priori* justification and *a priori* knowledge. A fallibilist can only talk confidently of *a priori* justification (meaning something like a non-empirical means of reliable truth-tracking, of reliably critically evaluating the credibility of certain beliefs). No claim to *a priori* knowledge could be conclusively justified, for a fallibilist.

This question of fallibilism will recur throughout Parts III and IV, and will be tied off in §7.5. This is an important aspect of our understanding of *a priori*, in the wake of the challenge of revisability and the externalist challenge. Even further, fallibilism in this sense may well be a non-obvious but inexorable consequence of the move from an acquaintance-based to an understanding-based approach to *a priori*.

(iv) Questions regarding relations between *a priori* and privileged access also surfaced above in the course of our discussions of ‘experience,’ and of

‘generality.’ Introspection and self-knowledge do instance a certain kind of immunity to error, and so have found themselves in the mix, when it comes to discussions of *a priori*. I agree with Burge (2000) (as discussed above) that the canon has not been completely consistent, when it comes to similarities and differences between what we might call introspection and rational intuition. In Part IV (§7.4), I will divide up the two pertinent cases and treat them separately, when it comes to the relations between introspection and *a priori*.

(v) Finally, then, for the relations between *a priori* and revisability. The challenge of revisability was one of the great shocks to the modal world-order in the twentieth century; and, looking ahead, incorporating this challenge into a satisfactory theory is a principle aim of this current project. Clearly, there are cases where what was once thought *a priori* is now thought false. For example, that space is Euclidean, that whales are fish, that there can be no such thing as a sub-atomic particle (since ‘smallest, indivisible’ was originally part of the sense of ‘atom’), etc., were all once considered to be justified *a priori* (and, arguably, analytically true), but are no longer so-classified; and it is hard to see how to conclusively rule out such a change in status of our beliefs.

On the view to be developed herein, the frameworks composed by our meanings, and which constitute our theoretical attempts to understand, are organic entities which change over time. Since *a priori* should be understood as framework-relative, what is justifiable *a priori* will change according to place and time. Nothing is absolutely and unqualifiedly *a priori*, but rather only *a priori* relative to certain axioms, postulates, presumptions. *A priori* should be seen as in large measure a matter of status, not just of content. There will still be a sense in which *a priori* justification is immune to counterexample; though—given the prevalence of conceptual evolution—that sense is importantly qualified.

§4.3: INTERIM CONSIDERATION OF SOME CANDIDATES

The next order of business is to discuss some exemplary instances from some major prospective categories of *a priori* knowledge. I will begin with considering some specific putative varieties of the *a priori* species, and then close this section with some general, interim remarks about relations between *a priori* and analyticity, as well as necessity. Note though I do little more than flag questions about certain candidates, at this stage. A conclusive discussion of these matters, on a constitutive *a priori* view, will have to wait until Part IV. I will begin by discussing some of the strongest, least contentious candidates, and move on to more contentious terrain.

It was more or less implicitly assumed, throughout the canon, that all and only necessary truths are knowable *a priori* and/or analytic; and so the cases with which we begin are strong contenders for all three categories. Paradigm examples of the necessary analytic *a priori* come from logic and mathematics:

1. $2 + 2 = 4$.
2. Squares have four sides.
3. No proposition is simultaneously both true and false.

These seem to bear the traditional rationalists' hallmarks of universality and generality. They are as strong a candidate as any for the status of being indifferent to anyone's noticing or grasping them, for not having just become true at any particular point. Bearing in mind the important distinctions between propositions expressed and the sentences that express them, these cases are, it seems, fashioned from the hardest of steel and immune to counterexample in metaphysical, semantic, and epistemic senses.

(Even these are not completely unanimous, as paraconsistent logics reject [3], and certain varieties of constructivism may balk at the claim that there is anything metaphysically necessary about [1] or [2]. Still, the claim that [1]–[3] are necessary, analytic, and *a priori* is rather orthodox, probably as close as one can get to unanimity in philosophy.)

Many hold that some moral judgments are also necessary analytic *a priori*:

4. Murder is wrong.
5. One ought to keep one's promises.

While this view is widespread, and a contributing factor toward many moves in metaphysics, moral philosophy, and theology, it is controversial. Many take the cultural variability of moral judgments to undercut any claim to necessity. The deeper problem here is the intuition of mind-dependence, that is, the notion that moral judgments, perhaps even more so than logical or mathematical ones, are contingent on the sorts of organisms that we are, or the ways in which we are encultured. Thus, while the view that the likes of [4] and [5] are necessary analytic *a priori* true has had lots and lots of defenders, this view is decidedly more controversial than the view that [1]–[3] are of that status.¹⁰

Truths by definition also seem to exhibit the relevant sort of metaphysical, semantic, and epistemic immunity to counterexample. That is, there is reason to put the following in exactly the same category as [2]:

6. Bachelors are unmarried men.
7. Vixens are female foxes.

To be sure, the careful work on conventions done in §§2.1, 2.2., and 3.2 is crucial here. There is nothing necessary about relations between bits of language and their meanings; but, once those conventions are set, it yet might be that what is semantically expressed by certain bits of language in a certain arrangement is necessarily so. Even given that, though, [6] and [7] seem more clearly and obviously analytic than necessary, *per se*. While there is bound to be large overlap between necessity and analyticity, [2] hooks onto a mind- and language-independent fact of the matter, in a way that [6] and [7] do not seem to. Hence, the inclination to take [6]–[7] to be entirely conventional, and so fundamentally analytic. In turn, the reasons to think that they are analytic are also, and thereby reasons to hold that their justification is, non-empirical.

[§]

We have unearthed some considerable reasons to expect these concepts to fail to be co-extensive, which reasons will be extensively developed in Parts III and IV. For now, here are some important cases which have been offered to show that necessary, analytic, and *a priori* are not co-extensive. Kant famously argued that there are synthetic *a priori* truths, such as:

8. Every event has a cause
9. $a^2 + b^2 = c^2$ (the Pythagorean theorem)
10. Humans are free agents who are subject to moral laws

Given Kant's assumption that necessity is one of the conditions for *a priori*,¹¹ these are quite plausibly necessary truths. However, he argues that they are not analytic, not true by definition. According to Kant, it is not a contradiction to judge that they are false (as it would be for the cases of [2], [6], or [7]).

Kant's ideas here are epochal and seminal. This was one of the most significant shakeups of the world-order, concerning *a priori*, in the entire history of Western philosophy. As mentioned above at §1.4, to this day, many proponents (e.g., BonJour [1998]) and opponents (e.g., Aune [2008]) of rationalism alike hold that rationalism is an interesting and significant position only if Kant is right that there are synthetic *a priori* judgments. Otherwise why fuss about rational intuition, if it is unable to amplify and add to the store of human knowledge?¹²

Much more recently, Kripke argues that there are necessary *a posteriori* truths, such as the following:

11. Heat is molecular motion.
12. Gold is the element with atomic number 79.

If science is in the business of discovering the essence of natural phenomena, then it seems that science aims to discover necessities. It is plausible to think that the likes of [11] or [12] are good candidates. For example, anytime molecules get agitated, it has to have the effect of increasing heat; and anytime heat is increased, that can only be because molecules are being agitated. However, despite the plausibility of the claims that [11] or [12] are necessary, there need not be anything *a priori* or analytic about them. Necessity, as Kripke (1972: 36) puts it, ‘in and of itself has nothing to do with anyone’s knowledge of anything’.

(It will be important to bear in mind below that the Kripke cases are always hypothetical, always of the form: IF science is right that *e* is the essence of P, then it is necessary but *a posteriori* that P is *e*. As a philosopher, not a scientist, Kripke is not in the business of determining the mind-independent nature of physical phenomena, but rather in the business of properly cataloguing the logic and semantics of these, and other, sorts of propositions.¹³ As we will see, this consideration complicates the question of the proper classification of these Kripkean examples of putative necessary *a posteriori* truths.)

Kripke also argues that there are contingent *a priori* statements. I will steer clear of many of his candidate examples (e.g., ‘metre,’ ‘Neptune’), on the grounds that some of them are attendant upon contentious theses in the philosophy of language whose critical evaluation lies beyond the scope of this project.¹⁴ However, the Kripkean or Kripke-inspired indexical cases of the contingent *a priori* (e.g., ‘I am here now’) do play an important role in the maps of the terrain developed in Part IV. Not only do they strongly suggest that analyticity or *a priority* are not sufficient for necessity; further, they prompt important reflections on the relations between semantic and epistemic immunity to counterexample.

[§]

Now to generalities. First, then, as to the relations between *a priority* and the analytic/synthetic distinction. While both of these notions have been taken by some to be suspect, anyone who concedes the intelligibility of both notions is bound to concede that there is overlap between *a priori* knowledge and analytic truth. Basic truths of mathematics and logic provide plausible examples, as do certain relational maxims (‘If X is north of Y, then Y is not north of X’). It is the worth of analytic *a priori* knowledge, when it comes to substantive problems in metaphysics or epistemology, which is contentious; the claim that (at least) some analytic truths are knowable *a priori* is, comparatively, fairly secure. (This becomes all the more plausible, less obscure, on the move from acquaintance-based to understanding-based accounts of *a priority*.)

As for whether all, or only, analytic truths are knowable *a priori*, that will take some digging and sculpting; how that matter stands on a constitutive *a priori* approach will be conclusively sorted out in Part IV. Kant famously argued that there are synthetic *a priori* truths, but many problems with Kant's arguments have been pressed in the intervening centuries. Ultimately, I will argue that all analytic truths are knowable *a priori*, but that there are plausible instances of the synthetic *a priori* (here indexical cases, such as 'I exist,' loom large).

Next then for the relations between *a priority* and the necessary/contingent distinction. Again, it is hardly objectionable that, provided that both necessary truth and *a priori* knowledge are coherent, then at least some necessary truths are knowable *a priori*. Again, elementary truths of mathematics and logic provide plausible candidates. However, while not staking any claims about all of Kripke's putative examples of necessary *a posteriori* truths and contingent *a priori* statements, ultimately, I will explain why I think that Kripke is right on both counts.¹⁵ This split between metaphysical necessity and any notions of epistemological modality is one of the deep and enduring legacies of the externalist challenge.

§4.4: SKEPTICISM ABOUT A PRIORITY

There is, on the one hand, pressure pushing this current section to be shorter than its closest predecessor §3.3, since some of the reasons for skepticism about necessity and analyticity carry over to *a priority* as well; so, those considerations are already out on the table, and do not need to be explicated afresh again here. On the other hand, though, in some respects *a priority* is the most objectionable of the three core modal notions, and so there is also pressure for this section to outstrip its parallel predecessor. Are we not finite and fallible epistemic agents? So, how could there be epistemic immunity to counterexample? Truth in virtue of meaning, maybe, and mind-independent facts, sure, but here we are talking about finite and fallible agents here.

For example, the notion of metaphysical necessity, of something which cannot be otherwise, seems to be perfectly coherent, even though there it may well be impossible to conclusively evaluate any particular candidate for metaphysical necessity. Similarly, most would concede the coherence of analytic truth; though many would question its usefulness or worth when it comes to heavy lifting in philosophy. In contrast, there can seem to be something deeply suspicious about *a priority*, and that per se. How can there be a source of justification—a way of reliably tracking the truth of mind-independent matters—over and above and apart from the

scientifically tractable causal relations between agents and their environments? *A priori* has seemed to many to be obscure, and perhaps even completely antithetical to a scientific world-view. In short, there seems to be something intrinsically supernatural about a non-empirical means of truth-tracking, and hence, despite its impressive historical pedigree, its very intelligibility is often questioned.

Following Peacocke & Boghossian (2000: 6), I will distinguish and briefly investigate three related sources of skepticism about *a priori*. Variants of each of them occur in the work of Quine; but the underlying sentiments pretty much thread through the history of the empiricists' opposition to rational intuition.

The first objection has it that the correct account of the growth of scientific knowledge suffices to refute the idea that an adequate epistemology requires the positing of *a priori* justification. The idea has it that science has no need of rational intuition, and science is the most effective path to knowledge yet developed. If we don't need rational intuition to explain scientific progress, and if scientific knowledge is the pinnacle of human knowledge, then it is just old-fashioned mystery-mongering for philosophers to insist on the need for the *a priori*. Maddy (2000, 2007) and Devitt (2005, 2011) provide recent statements of this line of objection; and both are representative of the prevalent naturalistic idea that the correct account of the growth of human knowledge is a Quine-inspired holistic web of belief.

The second, related objection to the intelligibility of the *a priori* is the obscurity objection: namely, there can be no satisfactory account of this mysterious, supernatural non-empirical source of justification. It could not possibly be fitted into a seamless scientific world-view, be reduced to the forces which we have good reason to believe govern the rest of the natural realm. In short, there can be no satisfactory explanation for how it is that some things could be non-empirically known.

Finally, the third line of objection is that *a priori* essentially entails certain instances of knowledge having a property which no instance of human knowledge could possibly have—such as infallibility, unrevisability, immunity to counterexample. Since the characteristic marks of the very idea of *a priori* are impossible, unsatisfiable, it follows that there can be no *a priori* knowledge for we finite, fallible agents. (See Kitcher [2000] for a recent example of such an argument.)

To take these three related objections in turn, then. The first is terribly presumptuous, assuming as it does that we have a satisfactory, comprehensive account of the growth of human knowledge in which the *a priori* plays no role. To the contrary, the idea that not even the actual history of science can be accounted for without appeal to *a priori* justification has

considerable defenders. (Cf., e.g., Pap [1946], Friedman [1992, 2000, 2011], DiSalle [2002], Richardson [2002], and Stump [2003, 2011, 2015], and Part III below for further discussion.) Devitt's (2011) efforts on this front are valiant, but hardly knock-down—his parting comment that 'many will remain unconvinced of the possibility of an empirical justification for the [contentious cases of] knowledge' (2011: 21) is an understatement. Maddy (2000) seems to admit the *a priori* in the back door after making such a show of chasing it out the front door (I will explain this allegation in §6.3). So, while there is something to this first line of thought, at this point in time it hardly supports categorical skepticism about the *a priori*.

Insofar as the second line of argument is merely that there is as of yet no satisfactory account of *a priori*, then that is hard to quibble with. Clearly, many hard questions remain, for proponents of *a priori*. However, insofar as this is supposed to support the conclusion that there is no *a priori*, then that is about as compelling as an argument that humans will never cure cancer, or travel to Mars, just because they have not yet managed to do so. True, in the case of the *a priori*, opponents point out that rationalists have had this problem in their laps for a millennia, and there are grounds for skepticism as to whether they have even made any progress. However, the constitutive *a priori* orientation developed herein takes a lot of progress to have been made on this question, first by Kant, and subsequently by several others.

As for the third line of argument, this is precisely the main order of business of this entire project. Precisely what do these challenges (reviseability, externalism) show about *a priori*? Groundwork for answering this objection has been laid throughout Parts I and II, and development of the answer will continue throughout Parts III and IV. The proof will be in the pudding.

[§]

One thing I should note, by way of segue, is that it is instructive to compare the force of these skeptical anti-*a priori* arguments along the dimension of acquaintance-based versus understanding-based orientations toward *a priori*. Not only the general ancient obscurity objection, but all of these more specific charges, apply less forcefully to the semantic intuition, understanding-based accounts than they do to the more traditional acquaintance-based approaches. The constitutive *a priori* view, which will be unpacked next, is, tellingly, a development within this less obscure, less contentious, semantic-understanding-based tradition.

As in the cases of necessity and analyticity, it is awfully hard to ground a firm, clear *a priorila posteriori* distinction, even though there clearly are paradigm cases at the poles (e.g., 'there is an apple pie baking nearby,' based on current olfactory experience, vs. 'all squares are four-sided,' based on a grasp of the concept 'square'). It may be best to say that there is a range

of shades of grey at the *a priori*-empirical divide. Conceptual evolution is a messy business, as concepts and meanings, like the languages they compose, are organic entities which change over time. However, again, the fact that there are shades of grey hardly entails that nothing is either black or white. The brute data of the poverty of experience remains. We are not about to be black-swanned by three- or five-sided squares, and that is not because we have already experienced all possible squares.

Moving on to Part III, then: now, to take this conceptual apparatus and apply it to the matter of exactly how some major twentieth-century developments in philosophy have affected the landscape of the ancient question of *a priori*.

NOTES

1. This distinction between the factive term '*a priori* knowledge' and the more cautious, purely epistemic term '*a priori* justification' is present in the literature under the guise of whether or not *a priori* includes or entails an indefeasibility requirement. Is it enough for a belief to be justified non-empirically for it to count as *a priori*, or is there a further condition along the lines of 'cannot be defeated by empirical evidence'? Kitcher (1983) and Field (2000) motivate an indefeasibility requirement, while Boghossian (1997), Peacocke (2000), and Casullo (2003) reject one. I am with the latter camp—as I have set things up here, indefeasibility is too much to ask of any purely epistemic notion. (It is no accident that Kitcher, for example, insists on an indefeasibility requirement on the path toward the larger dialectical game of establishing that mathematical knowledge falls short of *a priori*.) See also the discussion of 'fallibilism' about the *a priori* in the next section.

2. For example, one famous challenge stems from the Gettier (1963) cases of putative justified true beliefs which nonetheless fail to qualify as knowledge.

3. Herein lies another connection between *a priori* and innateness: poverty-of-experience arguments are also used to justify claims of innateness (as in the case of Chomsky [1967], for example).

4. Other famous elements of this camp of mathematician-Platonists also includes Descartes, Leibniz, Frege, and Godel (1944, 1947).

5. The allusion here is to Plato's *Symposium*, where Socrates reports that his insights came to him while lost in a trance, oblivious to what was going on around him. (Compare the passage from the *Phaedo* cited in §1.2.)

6. For explorations of distinctions between *a priori*, epistemic privilege, indubitability, and introspection, cf. the Introduction to Hanson & Hunter (1993).

7. Cf. Baehr (2003) for an argument in favour of the stronger claim that all of the various possible ways of drawing the divide face daunting problems. We will investigate in Part IV some self-knowledge cases in which difference senses of 'experience' results in different categorizations.

8. For discussion cf. BonJour (1998), Boghossian & Peacocke (2000), Casullo (2003), Schaffer and Veber (2011).

9. Burge (2000) investigates at length the relations between these points and the controversial roles which intuition plays in Kant's philosophy of mathematics, and also the ways in which the two notions of *a priori* are possibly conflated by Frege.

10. There is extensive discussion of a relevant distinction, between natural versus conventional reality, beginning at §5.4.

11. In agreement with much work in the rationalist tradition, Kant takes it as evident that at least some of our knowledge is universal and necessary, and argues that no empirical faculty could justify such claims. He explicitly assumes that anything known *a priori* is necessarily true.

12. This theme of the ampliative analytic recurs at various junctures, in various guises. Could semantic intuition amplify or add to the store of human knowledge? If not, how could it be of any relevance to Plato's problem, and if so, how could that be possible?

Indeed, many core, significant notions which pertain to this epistemology/language nexus—from Kant's (1781) synthetic *a priori* to Frege's (1892a) informative identity statements to Kripke's (1972) necessary *a posteriori* (and even on to Moore [1903] on the paradox of analysis)—are related to, if not guises of, these questions about the ampliative analytic.

The very idea of the ampliative analytic is also deeply affected by the move from acquaintance-based to understanding-based varieties of *a priori*—consider all the handwringing over 'from coherence to worth' in chapter 3, which gets more worrying to the extent that accounts of *a priori* get less obscure. Whether, in the final reckoning, my constitutive *a priori* view is a variety of moderate rationalism or of moderate empiricism hangs in the balance. (Cf. the final pages of chapter 6.)

13. Kripke (1972: 159): 'Certain statements ... if true at all must be necessarily true. One does know *a priori*, by philosophical analysis, that *if* such a statement is true then it is necessarily true. ... All the cases of the necessary *a posteriori* advocated in the text have the special character attributed to mathematical statements: philosophical analysis tells us that they cannot be contingently true, so any empirical knowledge of their truth is automatically empirical knowledge that they are necessary.'

14. In particular, careful examination of many of Kripke's examples would involve serious excavation of the semantics of proper names, descriptive names, natural kind terms, and theoretical terms more generally. I do get rather into natural kind terms in Part III, but otherwise do not make such serious forays into debates regarding the proper semantic treatment of these other sorts of terms herein.

15. To anticipate briefly: a case of knowledge counts as *a priori* if there is a non-empirical route to actual justification; necessity, or truth in all possible worlds, is light-years from here. These two properties might come apart most dramatically in indexical, *cogito* cases, or in cases of scientific discoveries of essence.

Part III

**ADAPTING AND APPLYING
THE CONSTITUTIVE A
PRIORI APPROACH**

Chapter 5

Two Major Challenges to the *A Priori*

This chapter speaks to two crucial ways in which developments in twentieth-century philosophy of language have altered the terrain, when it comes to *a priori* and the interrelations between metaphysics, semantics, and epistemology. We might think of these as the legacies of Quine and of Kripke, respectively—at the risk of doing a disservice to many of the others (cf. notes 1, 7) who did original and seminal work on these issues. They are the challenge of revisability and the externalist challenge, and they play a monumental role in my development of my preferred variant of the constitutive *a priori*.

The challenges are introduced and elaborated in §§5.1 and 5.3; and the work of absorbing the shock to the traditional world-order which they pose is begun in §§5.2 and 5.4, respectively, and further developed throughout the rest of Parts III and IV.

§5.1: QUINE AND THE CHALLENGE OF REVISABILITY

A longstanding, central plank of Quine's naturalistic campaign against immunity to counterexample is the idea that all human claims to knowledge are revisable.¹ Quine (1951: §6) is the classic statement of this viewpoint—containing such bold claims as 'no statement is immune to revision,' and that even revising the laws of logic would be '[no different] in principle [from] the shift whereby Kepler superseded Ptolemy, or Einstein Newton, or Darwin Aristotle'.

Looking back over the course of history, the claim that all beliefs are revisable looks to be fairly well supported. There is a large body of literature on the notion of a scientific revolution, and on the impacts of this concept within

our pertinent fields of epistemology and semantics.² For some examples, that space is Euclidean, that whales are fish, that there can be no such thing as a sub-atomic particle (since ‘smallest, indivisible’ was originally part of the meaning of ‘atom’), etc., were all once considered to be justified *a priori* (and analytically true), but are no longer so-classified; and it is hard to see how to conclusively rule out such a change in status of our beliefs. The question is what this challenge of revisability entails, for the core ancient notions of metaphysical, semantic, and epistemic immunity to counterexample.

For starters, I will briefly characterize three general lines of response to this challenge of revisability, which I will call ‘skepticism,’ ‘absolutism,’ and ‘revisionism.’ Modal skeptics follow Quine (1951) in holding that the challenge of revisability shows that analyticity, first and foremost, as well as necessity and *a priority* in its wake, is outmoded and untenable. The web of belief is seamless (to borrow Quine’s metaphor); all beliefs are subject to refutation and replacement. (Cf. Devitt [2011] for an illustrative recent statement of modal skepticism.) As we will see, the skeptic’s biggest problem is the original adequacy objection to radical empiricism—that is, the brute datum (of seeming immunity to counterexample) will not go away. (We are not about to be black-swanned by pentagonal squares, or by ever-childless grandmothers. It is a main aim of the work which follows to dig into the difference between atoms and whales, on the one hand, and squares and grandmothers on the other. About what, then, are we potentially black-swannable? Cf. especially §§5.4, 6.2, and 6.5.)

Modal absolutists dig in their heels and insist that immunity to recalcitrant experience is a central core component of an adequate epistemology; so to hold that (say) *a priority* is revisable is to change the subject. Frege (1884: 3) gives colorful expression to the absolutists’ creed: ‘An *a priori* error is thus as complete a nonsense as, say, a blue concept.’ Modal absolutism is a central plank in the traditional canon—for example, the view that knowledge of necessary truths can only be justified *a priori*, which is explicitly endorsed by Kant (1781: [B15]) among many others, depends upon this presumption that there is some supernatural potency about the *a priori*. However, absolutists must claim that these putative revisability-cases are actually cases in which one just mistakenly thought that one’s belief (e.g., about Euclidean geometry, atoms, or whales) was justified *a priori*. Among the problems with this option is that these notions (of *a priority*, analyticity, etc.) become only reliably useful for infallible agents, because agents like us could never conclusively establish the claim that something is justified *a priori* or analytically true. Thus, the absolutists’ notions of *a priority* and analyticity would be ill-suited to much work in epistemology or semantics. In any case, most contemporary theorists seem to be wary of modal absolutism—for example, it is explicitly considered and rejected by BonJour (1998:

Ch. 4), Field (2000), Peacocke & Boghossian (2000), Railton (2000), and Casullo (2003: Ch. 2). In short, even if the challenge of revisability does not suffice to support modal skepticism, it does amount to a rather strong case against modal absolutism. (Cf. the discussions of fallibilism and the defeasibility requirement, with respect to *a priori*, in §§4.1–2.)

Modal revisionists side with Pap (1946), Carnap (1950) and others against both the skeptics and the absolutists, in retaining immunity to counterexample but admitting that it is in some sense framework-relative. According to both Coffa (1991: Ch.10) and Friedman (2000: 370), the first clear articulation of revisionism occurs in Reichenbach (1920). Reichenbach alleges that Kant uses ‘*a priori*’ in two distinct senses—on the one hand, to mean necessary and eternal, and on the other hand, to mean constitutive of the concept of the object of knowledge—and goes on to argue that a moral of the theory of relativity is that the former be dropped while the latter retained. Revisionism seeks to define a principled middle ground between absolutism and skepticism about *a priori* justification, based on this notion of the constitutive but non-absolute *a priori*. The revisionists’ response to the challenge of revisability (for the core case of epistemic impunity) is to retain the concept of *a priori*, in many central senses of the term, while explicitly rejecting certain other of its traditional associations (such as necessity or infallibility). In addition to Carnap’s (1950) linguistic frameworks, this tack on the *a priori* is also widely associated with Wittgenstein (1921, 1953, 1969). Much interesting and challenging recent work on the *a priori* consists of variations on this revisionist theme—cf., for example, Friedman (2000, 2011), Railton (2000, 2003), Stump (2003, 2015).³

One important point to register at this stage is that my orientation on these issues is decidedly revisionist. The constitutive *a priori* view developed here is explicitly fashioned as a development within that tradition. Another is that, given the crucial differences between metaphysical necessity, analytic truth, and *a priori* knowledge, we should be open to the possibility that the appropriate responses to the challenge of revisability are relevantly different, among these three cases. This is one of many (connected) respects in which a gulf opens up between metaphysical modality on the one hand, and semantic and epistemic cases on the other. In particular, a sophisticated grasp of the ways in which analytic truth and *a priori* knowledge should be understood as revisable will be crucial, when it comes to refining our understanding of such notions. In contrast, metaphysical necessity is coldly indifferent to the challenge of revisability. (This particular motif has come up in both Parts I and II, and will be further developed as our story proceeds.)

Now to tie absolutism, skepticism, and revisionism to Plato’s problem (i.e., that human experience is finite and limited, and yet we seem to attain some knowledge which is universal and general). Absolutism strikes me as

singing loudly and proudly on the decks of the sinking Titanic, on this front. (I admire their audacity, but cannot square with their commitments.) Its current unpopularity is well-deserved. Skepticism, on the other hand, throws out the baby with the bathwater. Sure, things are difficult, when it comes to hard principled work on immunity to counterexample; but there is still the noble goal of an adequate epistemology, and the brute data, to be accommodated. Wouldn't it be nice if there were a viable *tertium quid*, between absolutism and skepticism, on the challenge of revisability? Is there a better alternative in the face of Plato's problem, than the absolutists' mere declaration of victory, and the skeptics' concession of defeat?

The aim of the rest of the book is to build exactly that, out of the ingredients cultivated in Parts I and II.

[§]

I am inclined to concede to the skeptics that the challenge of revisability shows up the untenability of modal absolutism. However, the move from here to a naturalism or skepticism pays an important price, and a price which need not be paid at that. For even if the challenge of reviseability amounts to a considerable case against absolutism, it does not amount to nearly as strong a case in favor of skepticism. What is this high price, exactly, then, and how would a non-absolutist go about not paying it?

The high price is, essentially, the ancient adequacy objection to empiricism, or all of the considerable reasons to posit *a priority* in the first place. The intuitive, principled difference between 'all swans are black' and 'all squares are four-sided' must be foregone. In other words, the price is precisely the notion of immunity to counterexample.

At a bit more length, modal skeptics must forgo belief in a privileged, constitutive connection between *understanding* and *justification*. To illustrate, compare the following pair (which we contrasted back in §4.1):

- [1] Squares have four sides.
- [2] Neptune has four moons.

For both cases, grasp of the meanings of the constituent bits affords an understanding of what would have to be the case for the sentence to express a truth. However, for the case of [1], grasp of the meanings also and thereby *justifies* the belief that what it expresses is true. Not so for [2], in which case understanding it does not come remotely close to providing justification for believing that what it expresses is true. Even though I know exactly what [2] means, I have no idea as to whether or not it is true; whereas it is far from clear that a correlative claim could coherently be made about [1]. Hence, [1] is a candidate example of this privileged connection between *understanding*

and *justification*: To understand [1] is *thereby* to be justified in believing it to be true.

This connection between understanding and justification is quite central to philosophy, both historically and conceptually—and especially to those orientations which take an understanding-based approach to *a priori*. Indeed, on some conceptions of the discipline, it is the very essence of philosophy as distinct from other theoretical enterprises; and so, for example, Anselm’s ontological argument, Descartes’ *cogito*, and Kant’s synthetic *a priori* are all instances of, or variants on, this general strategy of yielding justification from understanding. (However, the [UJ] connection is perhaps most strongly evident in the case of logical truths. This will be extensively discussed below §6.2.) Hence this price that the skeptic pays is quite high, with repercussions rippling from this corner of epistemology right through to conceptions of the methodology and subject matter of the discipline itself, as a whole.

Apart from this ‘threat of drastic shock’ sort of consideration about rejecting [UJ] connections, what are some positive reasons for keeping it around? What can it buy us?

A lot: these [UJ] connections lie at the heart of any understanding-based approach to *a priori* (of which the constitutive *a priori* orientation is an exemplary instance). They are the basis of the link between analyticity and *a priori*—given that understanding can (in some cases, the boundaries of which are to be limned below) suffice for justification, the resulting set of analytic truths are justified *a priori*. If there are [UJ] connections, then analyticity can ground (at least some cases of) *a priori* justification. I’ll call this [UJ] principle 1:

$$UJPI: [UJ \rightarrow (A \rightarrow AP)]$$

UJPI will be extensively developed throughout Part III, and will play a role in the maps of the terrain detailed in Part IV. Before that, though, much work needs to be done in terms of showing how the [UJ] connections are to survive the challenges of revisability and externalism.

Wherever the relation between semantic intuition and rational intuition has popped up, these [UJ] connections are what we are stalking. These are the cases in which semantic intuition can give all that was wanted, in terms of epistemic justification, from rational intuition. These are the cases in which semantic immunity to counterexample suffices for epistemic immunity to counterexample. So, this is the realm of the un-black-swannable, whose precise contours it is a primary aim of Parts III–IV to chart. The realm of the un-black-swannable is the range of cases in which

semantic intuition can underwrite the [UJ] connections, and hence [UJ] principle 2:

$$UJP2: [UBS \leftrightarrow (SI \rightarrow UJ)]$$

So, arguably, nothing less than a non-obscure but adequate epistemology—between absolutism and skepticism—is at stake here.

[§]

The basis of the challenge of revisability, then, is that U (i.e., understanding) seems to be too variable, across times and places, to cement any such firm J (justificatory) foundation. Understanding seems to be relative to various contingent factors. Precisely how the challenge of revisability is to be absorbed, on my orientation, is developed first in the next section, and then charted in the next two chapters. This kind of relativity (i.e., in deep tension with absolutism but a far cry from skepticism) became clearly evident by mid-twentieth century; the varieties of revisionism have blossomed since.

To sum up: Given that modal absolutism is off the table, and that modal skepticism is unwarranted and inadequate, revisionism is a worthwhile research project, and the constitutive *a priori* is a promising variety of revisionism. The [UJ] connections are worth keeping around for a host of reasons, including especially their promise to ground a non-obscure but adequate stance regarding Plato's problem.

§5.2: FRAMEWORKS AND THE CONSTITUTIVE *A PRIORI*

The challenge of revisability is that as we humans are fallible and limited, all claims to knowledge seem to be subject to revision over time. Absolutists dismiss the challenge by insisting that such modal notions as analytic truth and *a priori* knowledge are essentially unrevisable. Skeptics, in contrast, hold that the challenge of revisability shows up as mythical these historically significant ideas of immunity to counterexample. We are now engaged in the task of developing a revisionist response to the challenge of revisability. In particular, this present section will continue the development of my favored response to this challenge, which was described in a preliminary way above in §5.1, and will be developed more extensively below—cf. especially throughout chapter 6.

The core revisionists' idea is Kantian in spirit, though it self-consciously departs from some elements of Kant's view. Revisionists seek to define a principled middle ground between absolutism and skepticism, based on this

notion of the constitutive but non-absolute *a priori*. Hence, the revisionists' response to the challenge of revisability (for the core case of epistemic modality) is to retain the concept of *a priority*, in many core senses of the term, while explicitly rejecting certain other of its traditional associations (such as necessity or infallibility). There is no entailment from 'constitutive *a priori*' to eternally or necessarily true, though (as we will see) there do remain some clear senses of immunity to counterexample—once the challenge of revisability is properly digested.

For a revisionist, *a priority* (as well as, relatedly, analyticity) must be understood as relative, in a certain sense—for example, to a linguistic framework for Carnap (1937, 1950), to a language game or world picture for Wittgenstein (1953, 1969), to a theoretical framework in Friedman's (1992, 2000, 2011) distinctively Kantian take on this same core idea. (I will stick with Carnap's familiar vocabulary and use the term 'frameworks.' My usage is general, such that distinct language games, pictures, theories, etc., constitute different frameworks.) However, this relativity stops well short of modal skepticism (i.e., dismissing the very idea of immunity to counterexample as folly). Such special modal notions as *a priority* must be understood not as marking off some queer kinds of objects of knowledge, but rather as indicating a special status attached to certain beliefs. To call something *a priori* is, in part, to say something about the role which it plays in the relevant framework.

Many traditional approaches to *a priority* regard *a priori* knowledge as essentially involving a special sort of content (i.e., self-evident grasp of bulletproof superfacts, which glow with luminous certainty). However, proponents of the contingent *a priori* take *a priority* to be also essentially a matter of status, not just of content. *A priority* must be understood not as marking off some special kinds of objects of knowledge, but rather as indicating a special role, function, or status attached to certain tenets. To call something *a priori* is to say something about the role which it plays in the relevant framework. The *a posteriori* beliefs are those that the agent treats as being subject to the tribunal of experience; the *a priori* beliefs are subject to a higher court.

Consider, for example, an agent who sincerely avows the universal generalization that *every event has a cause*. (This is roughly based on an example discussed by Railton [2000: 178].) Further questions might arise as to the precise content and status of this belief—for example, is this a regulative rule for the agent, the so-called 'principle of sufficient reason' (i.e., any conceivable event *must* have a sufficient cause), or is it rather an inductive generalization (i.e., as far as I know, every event observed to date by any credible observer has had a sufficient cause)? One way to tell is to present the agent with a putative counterexample; say, an alleged uncaused event in the quantum void. To the extent that the agent responds with categorical denial—presuming that there has to be a sufficient cause there, whether or not anyone has yet detected

it—that indicates that this particular belief is an *a priori* regulative rule. If, in contrast, the agent is willing to defer to scientific experts on the matter, and to withdraw or qualify the original universal generalization, then that shows that it was all along an *a posteriori* inductive generalization. Thus, ‘*a priori*’ does not simply apply to the content of a belief, but, rather, also has essentially to do with its status, or its place in the relevant, operative framework.

To cite a couple of examples from Carnap (1950), that there are numbers is a constitutive *a priori* principle of the framework of elementary arithmetic, and that there are ordinary physical objects is a constitutive *a priori* principle of the framework of folk physics. Considered internally, from within the frameworks, such principles have the status of immunity to counterexample—they are treated as simply not being subject to empirical disconfirmation. They are rather constituent elements of the rules of the game, without which various pertinent sorts of questions could not be asked, or conjectures could not be tested.⁴ Carnap proceeds from this point to dismiss many traditional philosophical questions—for example, ‘Yes, but do numbers, or physical objects, really exist mind-independently?’—as mistaken pseudo-questions, confections of the crucial distinction between internal questions (within the framework) and external questions (about the framework). However, while modal revisionism is essentially allied with this Carnapian (neo-Kantian) notion of the framework-relative constitutive *a priori*, it need not take on any such positivistic meta-philosophical theses.

And note well the clear sense of revisability here. Framework-relative *a priori* does not involve Platonic, supernatural grasp of luminously certain, eternal superfacts. The frameworks of mathematics and of folk physics do evolve over time, with the attendant corollary that which principles get treated as having this status of immunity to counterexample—as conditions for the possibility of asking clear questions and posing testable hypothesis—also change over time. For example, it was once justified *a priori* that negative numbers have no square root (for empirical investigation was not required for, or even relevant to, establishing that the product of no number times itself could be a negative number). However, our current framework now includes as an axiom that the imaginary number i is the square root of -1 . And note (cf. note 4) that the introduction of i is not smoothly analogous to the discovery of another previously unknown moon orbiting Neptune. It rather constituted a change in the rules of the game.

Further illustrations of the modal revisionists’ notion of a constitutive *a priori*, as well as arguments in favor of the indispensability of the notion in accounting for scientific progress (contra a prevalent brand of holistic, Quinean, anti-*a priori* naturalism) are assembled by Pap (1946, 1958), Friedman (1992, 2000, 2001) and Stump (2003, 2011, 2015), among others.⁵ Pap’s (1946) driving idea is that every scientific theory is built on fundamental

principles which must be treated as unassailable for the purposes of framing hypotheses, but nonetheless which principles are treated as having such a status (can, should, and do) change over time. Friedman (2000) develops some specific examples along this vein in considerable detail. For example, in the case of Newton's scientific advances, certain principles underlying the calculus have to be treated as *a priori* in order to even formulate, let alone test, Newton's laws of motion; and, in turn, Newton's theory of gravitation could not even be intelligibly formulated without taking the laws of motion as *a priori*, as not subject to empirical disconfirmation. Friedman (2000: 377) describes the epistemological upshot of such episodes in the history of scientific progress thus:

What characterizes the distinguished elements of our theories is ... their special constitutive function: the function of making the precise mathematical formulation and empirical application of the theories in question first possible.

This is the heart of the notion of *a priority*, for a modal revisionist. Beliefs that have the status of *a priority* are those which play a certain kind of structuring, regulative role in the framework. They are those that are treated by the relevant agent (or community) as being immune to empirical counterexample.

This move from taking *a priority* to subcategorize the alleged special kind of content of certain special propositions to taking it to be, at least in part, a question of status, function, or role within a framework, is one core pillar shared by all constitutive *a priori* views. Another is the notion that something must be treated as having this special status in order to get any intelligible inquiry up and running. In order to clearly pose precise, intelligible questions—to isolate variables for testing—some things need to be treated as immune to empirical counterexample. (Examples will be developed below, especially at §§6.1–3, to illustrate some of these generalities.) Such things are thereby constitutive of the framework of inquiry. Frameworks can and do change over time; and there are mundane cases in which multiple distinct frameworks are brought to bear on a particular problem. Nonetheless, in any specific case, there have to be some things with this status of immunity to counterexample to constitute a framework which can then be applied in inquiry.

So, whereas the absolutist treats *a priori* knowledge as essentially involving a special sort of content (i.e., self-evident grasp of superfacts, which glow with luminous certainty), for the revisionist *a priority* is also a matter of status, not just of content. *A priori* knowledge is revisable, on this orientation; though to make such a revision is a much more drastic matter than revising beliefs that lack this status. To revise the *a priori* is to change the framework (or language game, theory, etc.).

The exact nature of this constitutive *a priori* varies among the varieties of revisionism—for the Wittgensteinians, the crucial distinction is that between

rules and propositions (i.e., between the rules of the game and the moves which may be made according to those rules); for the Carnapians, the key distinction is between the pragmatic and conventional criteria which define a framework and the things which then become sayable or decidable within that framework; etc. (Much more on this in §6.1.) And note well that this is not merely a bifurcated, twofold distinction. Frameworks are often re-evaluated, revised and updated, in more or less drastic ways; we often encounter complex situations to which multiple distinct frameworks may be simultaneously applicable, and the relations between these distinct frameworks can be multifaceted and dynamic. Conceptual evolution can only be framed as a neat narrative in hindsight.

It is a main order of business in chapter 6 to mine and sort these kinds of complexities. (See also §5.4 for discussion of some crucially different kinds of framework.) Across the diverse spectrum of distinct kinds of framework—from agriculture to cosmology, from meta-ethics to mathematics—the particular balance between internal and external questions may well vary widely, as may the sorts of considerations which matter when it comes to engaging with both sorts of questions.

Note also that, even despite this stress on status as opposed to just content, obviously not all contents are equally suited to such a status. For example, in §6.2 I will argue that $\&$ -elimination is much better suited for *a priori* than the law of excluded middle; and examples of contents mistakenly treated by others as being immune to counterexample are not hard to find (e.g., ‘White males are intellectually superior,’ ‘Bad things happen in threes,’ ‘All that happens is for the best, because it is God’s will’). The evaluation of frameworks as more or less reasonable, based on differences among the contents which are taken to be *a priori*, has clear appeal both within and beyond its promise to help make sense of the notion of scientific, political, and philosophical progress.

In any case, getting back to [UJ] connections, the core idea here is that *a priori* (which includes our focal notion of logical truth as a distinctive sub-case) should be understood not as marking off some queer kinds of objects of knowledge, but rather as indicating a special status attached to certain basic tenets. To call something *a priori* is to make a claim about the kind of basic, structuring, regulative role which it plays in the relevant framework. So note that this kind of revisionism in general is allied with an understanding-based, as opposed to an acquaintance-based, approach to *a priori*.

This notion of a framework, and its relevance to the challenge of revisability, among many other things, will be further elaborated below. My constitutive *a priori* variant of modal revisionism holds that there is immunity to counterexample but it is framework-relative. I will continue to build the case that that is the best way forward from this challenge, all things considered, for grounding a non-obscure but adequate answer to Plato’s problem.

§5.3: KRIPKE AND THE EXTERNALIST CHALLENGE

We begin our engagement with the externalist challenge with a quick trip back to some basic notions from Part I (especially §2.3), to set things up. For the externalist challenge affects precisely the relations between meaning and extension; which thus problematizes considerably the notions of semantic and epistemic immunity to counterexample, and their relations to metaphysical necessity. When it comes to the core [UJ] connections and principles, at the heart of any understanding-based approach to *a priority*: this potential blockage between meaning and extension threatens to undermine whether U (understanding) can supply anything remotely like J (justification); and so thereby whether semantic intuition can be relied on for anything substantive in epistemology.

Hence, in addition to the challenge of revisability, and dovetailing with it in many respects, the externalist challenge too threatens to undermine the core [UJ] connections. In general, Part III is all about saving [UJ]—and so what semantic intuition and understanding-based accounts of *a priority* can do for Plato's problem—from the two-pronged assault which I am associating with Quine and Kripke.

I will use the term 'traditional internalism' to designate a certain conception of the relation between meaning and extension, which went virtually unchallenged until well into the twentieth century. Canonical sources which state clear allegiance to this orthodoxy include Plato (1928b: 324A–343A) and Locke (1690: Bk 3, I–III). These presumptions are guiding principles in seminal work in the philosophy of language by Mill (1843), Frege (1892), and Russell (1918), and are explicitly defended as recently as Strawson (1959) and Searle (1969).

Details differ, but the general internalist picture is this: every term is semantically associated with a *meaning* which specifies the conditions for membership in the term's *extension*. Competence with a term is a matter of associating it with the appropriate meaning, which is made manifest by the agent's ability to distinguish the extension from the anti-extension (in normal contexts). On the traditional view, the criteria for the correct application of a term are introspectively available to competent agents.⁶ These criteria for correct application, and hence the content of the propositions expressed by their utterances, are completely transparent to individual agents—there is nothing hidden from view, no reason why we would have to invoke something external to an agent in order to individuate the contents expressed or entertained. Individual agents who are semantically competent are autonomous as to the conditions that determine the extensions of their terms.

Wittgenstein (1953) is one influential critic of this traditional internalism. He points out that while most of us are rather good at distinguishing

the extension from the anti-extension of the term ‘game,’ for example, we are rather horrible at articulating any meaning that specifies what all and only games have in common. Wittgenstein is also thoroughly critical of the presumption of first-person authority about meaning, insisting rather that the criteria for the correct application of terms crucially depends on the practices of a community. Strawson (1959) and Searle (1969) both attempt to accommodate some of Wittgenstein’s insights, within general internalist confines.

Another forceful challenge to the traditional orthodoxy comes in the 1970s (though it was certainly influenced by Wittgenstein, among others).⁷ Consider, for example, an agent who associates with the name ‘Columbus’ the inaccurate meaning ‘the first European to sail to North America,’ or who associates with the name ‘Einstein’ the vague meaning ‘a famous physicist.’ First, Kripke points out that these sorts of cases are fairly common, much more representative than the small handful of tendentious examples discussed within the traditional orthodox literature (e.g., ‘Bismarck’ means ‘the first Chancellor of the German empire’). Second, Kripke motivates the claim that such speakers nonetheless count as competent with these terms—they are able to participate in the interchange of information about Columbus and Einstein—despite not having any introspective grasp of the conditions for the term’s correct application. (The agent knows nothing to distinguish Einstein from Heisenberg or Feynman, and the condition associated with ‘Columbus’ probably picks out some ninth-century Viking.) Third, Kripke argues that this shows that whatever it is that determines the extension of a use of a term, it must be distinct from the often vague and shoddy information that constitutes the meaning which the speaker associates with it. In general, meaning (i.e., the information which the speaker associates with a term) need not determine extension. The conditions for the correct application of a term need not be accessible to competent speakers.

‘Externalism’ is an apt label for this line of thought in that the upshot seems to be that (at least in some cases) something external to the agent must be invoked in order to determine the extension of the terms they entertain and express. In addition to proper names, the externalist challenge also forcefully applies to natural kind terms. Competence with such terms does not depend on a grasp of the precise criteria for their correct application—even if I couldn’t tell whether some non-typical specimen is or is not a tiger, still I count as competent with the term ‘tiger.’ Prevalent candidates for semantically relevant external factors include: (i) the causal-historical chain of transmission of the expression tokened, (ii) facts about the actual nature of the ambient environment which may be inaccessible to ordinary speakers (e.g., differences between gold versus iron pyrites, H₂O versus XYZ), and (iii) the states and doings of certain specific sub-sets of the linguistic community in which

the speaker is immersed (such as Putnam's [1975] 'experts,' or Evans' [1982] 'producers').

[§]

I will use the term 'the externalist gap' [EG] to denote this phenomenon wherein meaning (i.e., the conditions which competent speakers associate with a term) seems to be distinct from what determines the term's extension. To say that there is an [EG], for a particular use of a particular expression, is to say (in Kripke's [1972] terms) that separate answers seem appropriate for what 'gives the meaning' versus what 'fixes the reference.' Alternatively, we could define the [EG] using Putnam's (1975: 225) "two [traditionally] unchallenged assumptions": (i) knowing the meaning of an expression is a matter of being in a certain intrinsic state, and (ii) the meaning of an expression determines its extension. The core of the externalist challenge to traditional internalism is the notion that, in general, nothing can do both jobs (i) and (ii).

Traditional internalism holds that one univocal 'meaning' (in a fairly clear sense of that notoriously ambiguous term) is both that the grasp of which constitutes competence, and that which determines extension. (For example, the meaning of 'triangle' is something like 'three-sided closed plane figure'; a grasp of that constitutes competence with the term, and anything that satisfies that counts as a triangle.) The externalist challenge to traditional internalism, then, has it that, in general, no one thing can play both of these roles. For any particular semantic property S: if S is intrinsic to individual speakers, then S does not in general determine reference; and if S determines reference, then S is not in general intrinsic to individual speakers. Putnam (1975: 249–50) puts this point by exploring the ways in which 'the traditional problem of meaning splits into two problems'—that is, 'determination of extension' versus 'describing individual competence'. A more Kripkean (1972) spin would be to criticize Fregeans for using the term 'sense' in two distinct senses, and to then develop the distinction between 'giving the meaning' and 'fixing the reference'.

Herein also lies the reference fixer—as distinct from both meaning and extension—discussed above at §§2.3 and 3.2. In cases in which there is an [EG] (i.e., the meaning is distinct from the criteria for membership in the extension) there arise questions as to the nature and workings of the reference fixer. This is a significant aspect of the externalist challenge, and it problematizes precisely the putative transparent relation between meaning and extension—and hence between the U (i.e., what constitutes understanding) and the J (i.e., whether it can amount to justification).

[§]

I will not go further into the exact substance of the externalist challenge to traditional internalism, but rather will take it as commonly known and relatively uncontentious. Along with most, I hold that Donnellan (1970), Kripke

(1972), Putnam (1975), Burge (1979), and others have developed a very strong case for an [EG], at least in some cases. (Further, I will take Putnam [1975] on the distinction between (i) and (ii), discussed above, and Kripke [1972] on ‘giving the meaning’ versus ‘fixing the reference,’ to be completely interchangeable ways of articulating this [EG].)

One significant way to divide up varieties of semantic externalism, then, concerns the range of cases which are amenable to this [EG]. To illustrate, a view toward the moderate end of the spectrum of options holds that there is only an [EG] in cases of deferential uses of proper names and natural kind terms (e.g., ‘*Feynman* is a physicist,’ ‘*Molybdenum* is a metal,’ said by speakers who would explicitly disavow any ability to distinguish their referents from other physicists or other metals). In contrast, a view toward the opposite extreme holds that the [EG] is applicable to all thought and talk—even to ordinary, competent speakers’ everyday usage of the simplest possible terms (e.g., ‘and,’ ‘hunter,’ ‘bachelor,’ ‘triangle’). In general, there are considerable reasons to be given in favor of both of these extremes, and there are also grounds for lots of principled intermediate views.

I will distinguish three different theses about the range of the [EG]—which I will call the Pragmatic, Semantic, and Metasemantic theses. The Pragmatic thesis takes the [EG] to be a property of certain limited type of linguistic *usage*; the Semantic thesis takes the [EG] to be a property of certain limited type of linguistic *expression*; and the Metasemantic thesis takes the [EG] to be a general discovery about the nature of *language*. By §6.5, I will have assembled the ingredients for an argument for the Pragmatic thesis, and, relatedly, in favor of a relatively moderate semantic externalism. One main aim of chapter 6 is to explain that moderate externalism is how an understanding-based, constitutive *a priori* approach best absorbs the externalist challenge.

So, to sum up: there is a very strong case for an [EG], at least in some cases, and it is certain to have considerable impact on our maps of the terrain at which epistemology overlaps with semantics and metaphysics. But the proper upshot of this externalist challenge is still somewhat up in the air. Some of the hard work to be done in order to sort this out includes: (a) getting a firm grip on the extent of the externalist challenge—that is, how far beyond the cases of proper names and natural kind terms do these arguments apply?; and (b) getting a better handle on the mechanisms that externalists hold to play a role in determining reference—and in particular on the notion of deference, which plays a critical role in virtually any post-Wittgensteinian theory of reference. Both of these matters are engaged in the considerable depth in chapter 6.

[§]

When it comes to the externalist challenge to traditional conceptions about meaning and extension, so-called ‘incomplete mastery’ cases are centrally

at issue—that is, speakers of whom there is considerable reason to count as competent with an expression, but who nonetheless are not able to articulate effective criteria for its correct application, which distinguish the extension from its complement. If incomplete mastery does not entail semantic incompetence, then the ability to fix the reference is not criterial for grasping the meaning. The Wittgensteinian line of thought that some or most of our concepts are not constituted by necessary or sufficient criteria poses one line of challenge to traditional internalism here. But the externalist challenge, while perhaps properly seen as a further step in this same direction, is a more categorical jolt. Regardless of what one thinks about exactly how to characterize some or most of our concepts, externalists argue that what matters for competence and what determines the reference are generally distinct sorts of question, the answers to which involve rather different sorts of considerations.

When it comes to natural kind terms, for which the case for a division of linguistic labor is especially strong, the intuition that incomplete mastery is compatible with competence is fairly robust. We are content to count Kripke as competent with ‘tiger’ even though he might not be able to say of some non-typical specimen whether or not it ought to be classified as among the tigers, and to count Putnam as competent with ‘gold’ even though he is potentially subject to dupe by clever counterfeit. Fair enough, so far. Of course, Putnam (1975: 233) himself points out that ‘some words do not exhibit any division of linguistic labour: “chair,” for example’. (As we will see in §6.5 below, there are passages in Kripke [1972, 1979] which suggest that his externalism is much more moderate than radical.⁸) Even the move to Putnam’s (1975) ‘beech’/‘elm’ case might be argued to be a difference of kind and not degree—perhaps our community demands more to count as competent with these terms, than merely that they name ‘some kind of tree’.

Things get shakier still with, say, Salmon’s (1989) claim that competence with the terms ‘catsup’ and ‘ketchup’ is compatible with believing that they name distinct condiments. (That is, exactly akin to a Hesperus/Phosphorus case, one could be competent with both ‘ketchup’ and ‘catsup’ while thinking that they name distinct condiments.) Really? *REALLY?* Competence? I take it that the obvious traditional internalist response—that is, to the contrary, such an agent ipso facto falls short of competence with at least one of these terms—clearly has considerable purchase in this case. It takes some audacity to try to twin-earth up an [EG] for ‘triangle’ (or ‘grandmother,’ ‘fortnight,’ or ‘and,’ to name a few).⁹

§§6.3–5 will consider these matters closely. For now it suffices to underline that incomplete mastery cases encapsulate what is at issue here, when it comes to differences between internalists and externalists. And so, again, when it comes to the question of the range of the [EG], is the moral that there exist certain distinctive, limited sort of cases in which competence is

compatible with incomplete mastery? (In Fregean terms first broached in note 15 of chapter 1, does sense *not always* determine reference, or rather *generally never* determine reference?) Or that competence is, quite generally, compatible with incomplete mastery, for any linguistic expression?

A related bone of contention between internalists and externalists concerns the transparency of meaning, or whether typical speakers individually have introspective access to the criteria for the correct application of the expressions with which they are competent. The transparency of meaning was an axiomatic presumption for traditional internalists—one gets a strong whiff of an axiomatic commitment to transparency in reading, say, Russell (1918) or Frege (1892a). (Transparency requires no defense by argument; it is analytically entailed by what ‘meaning’ means! Further, it is crucially presupposed in, for example, the characteristic distinctive inferences which Frege and Russell draw from the informativeness of statements of the form ‘ $a=b$ ’—in Frege’s case, to a difference in sense between ‘ a ’ and ‘ b ’; in Russell’s case, to the conclusion that at least one of ‘ a ’ and ‘ b ’ is a description in disguise.) Indeed, it is only once something like Kripke’s (1972) distinction between ‘giving the meaning’ and ‘fixing the reference’ is drawn that failures of transparency for competent speakers become a clearly intelligible possibility. However, most contemporary philosophers take the seminal externalist arguments to have established counterexamples to transparency. (Most but not all—even in the wake of the seminal externalist arguments, we still get Dummett’s [1978: 131] flat assertion that ‘transparency’ is ‘an undeniable feature of linguistic meaning’—the preservation of which intuition is one of the driving motivators for two-dimensional approaches to semantics [cf. Chalmers {2006}].)

Transparency is of course just a metaphor, but an apt one. Given that the meaning fixes the reference, then there is in principle nothing hidden from the view of competent speakers: the light of the individual competent speaker’s mind illumines the boundaries and contours of the extension. The criteria for the correct application of a term are introspectively available to competent agents, who are thus autonomous as to the conditions that determine the reference of their terms, according to this aspect of traditional internalism. One last corollary question, then, probing further into differences between the Pragmatic, Semantic, and Metasemantic theses about the [EG]: Have the seminal externalist arguments shown that the meanings of certain distinctive kinds of expression are not transparent? Or that linguistic meaning, in general, is not transparent?¹⁰

[§]

Hence, this externalist challenge is deeply relevant to several core questions about immunity to counterexample, and about the nexus at which the

philosophy of language overlaps with metaphysics and epistemology. It also ties in with the challenge of revisability in complex and interesting ways, further complicating the framework-relativity of the U (understanding) aspect of the core [UJ] links and principles. (Quine and Kripke in superposition!) The challenge of revisability strongly suggests that U is historical, perspectival, fleeting, and the externalist challenge threatens to have identified a blockage between U and J.

Revisability first: is anything really immune to counterexample? Here we will have to distinguish between the cases of metaphysical, semantic, and epistemic impunity. Working through this challenge will prompt refinement to their interrelations, to the relation of each to semantic intuition, and ultimately to Plato's problem.

As for externalism, this prompts much refinement to our understanding of the notions of meaning, extension, reference determiner. It also limits the metaphysical conclusions which could be supported by anything having solely to do with reflection on meaning, content. (Does, or can, grasp of meaning bring anything in its train that is of epistemic or metaphysical import?) This is particularly true of questions which pertain to analyticity and *a priori*—for those notions are, to a large extent, precisely about transparent access to a certain sort of truth or knowledge. The externalist challenge (to the connections between meanings and extensions) threatens to undermine that access. Immunity to what, and if so how?

The two challenges dovetail in that the gulf between analyticity and *a priori* on the one hand, and necessity on the other side which is required to answer revisability—which has been a recurring motif since the latter pages of chapter 1—is also a moral of the externalist challenge. Like revisability, externalism also affects discourse about mind matters and language matters of metaphysics differently than it affects discourse involving categories which humans construct rather than discover. We will now dive headlong into this chasm.

§5.4: NO STATEMENT IS TRUE BUT REALITY MAKES IT SO

Framework-relative modal revisionism will be my answer to the challenge of revisability, as developed on an understanding-based constitutive *a priori* orientation. What, similarly, answers the externalist challenge? One crucial part of the answer concerns principled differences between different kinds of frameworks, and the different mechanics appropriate across this divide, when it comes to relations between meanings and extensions.

We saw in §4.3 that Quine's flat-footed insistence that 'No statement is true but reality makes it so' has played a role in shaping various pictures of

modal space. This present section continues the work started there, in terms of honing and refining what that crude dictum might mean, and what consequences it should have. Perhaps the most important present point is that if this tenet is to be upheld, then it is crucial to distinguish between certain different kinds, or aspects, of reality. This is an important step in a satisfactory, refined response to the externalist challenge. Generally, this present section is going to unpack some material that chapter 6 is going to then spread out and develop.

A key point is that Quine's dictum is going to apply rather variously, across different kinds of frameworks. Consider for example the different ways in which 'No statement is true but reality makes it so' apply to, or impose constraints on, 'Aluminum is a metal' versus 'Widows are formerly married women whose spouse has died.' Only the former *targets* and makes a *judgment* about mind-independent reality; the latter does something more like *categorize* something which is *presupposed* rather than judged upon. Non-empirical access to the 'facts' or 'reality' at stake here are really quite drastically different matters (i.e., access to the categories themselves versus access to what is intended to be categorized). So, given the huge differences between the kinds of concept that 'aluminum' and 'widow' are, questions about *a priori* and the externalist challenge are going to apply rather diversely to them. (Compare the different senses of 'factual' discerned in §4.3.)

Conventionalism (i.e., roughly, human convention is the source of what otherwise might seem to be metaphysical necessity) may not be an adequate approach to all immunity to counterexample, but it still might be the right thing to say about some of it. Human cognitive activity is partly constitutive of reality when it comes to some of our thought and talk. And this will be crucial when it comes to not only the externalist challenge, but also Plato's problem, and staking out the un-black-swannable turf. (The deep and crucial [EG]—and, relatedly, the varieties of immunity to counterexample—is going to apply rather diversely across this divide among distinct sorts of framework.)

This divide is most directly and explicitly about how to absorb the externalist challenge, though the ways in which it is also pertinent to the challenge of revisability will also be charted. Since this fissure affects the notion of a framework, in fairly deep and far-reaching ways, and frameworks in turn are tied integrally to our responses to both of these challenges, many such considerations apply to both phenomena together.

Recall too the point made in §1.4 above, that one fairly canonical way to divide off the moderate rationalists from the moderate empiricists (who can look rather similar, as both camps agree that there is *a priori* justification but reject acquaintance-based accounts) concerns whether it is possible to attain *a priori* knowledge 'of the world.' Thus, empiricists would typically toe the

Humean line that *a priori* knowledge is limited in scope to our own ideas or concepts, while the rationalists would espouse a more ambitious and bold conception of the range of *a priori*. This fissure will be further developed, from this present section and on through what follows. The question of the ampliative potential of semantic intuition will be crucial in general, in developing my constitutive *a priori* account, and decisive in particular, when it comes to whether that account should be seen as a variety of rationalism or of empiricism.

[§]

To begin, let us distinguish between social-conventional reality and mind- and language-independent reality. (Cf. Cassam [2000] for discussion; he acknowledges Locke and Kant as guiding influences.) Cassam's (2000: 59) examples of social-conventional phenomena include that January has 31 days, and that suicide is the taking of one's life. We enjoy a sort of privileged access to social-conventional phenomena, precisely because what grounds these categories, and holds them in place, is our thought and talk. There is no mystery as to how statements about all Januarys having 31 days, or all suicides being the taking of one's own life, could be knowable *a priori* because we intentional agents are ourselves co-constructors of the data. There need be nothing remotely supernatural about it. (This might seem to run into tension with the widespread externalist tenet that meanings and concepts are fundamentally the property of a community, not an individual, but stay tuned for the discussion [in §§6.3–5] of transparency, deference, and the limits of semantic externalism.)

In contrast, natural kind terms are the paradigm case of expressions which concern mind- and language-independent phenomena. To use a term as a natural kind term involves a deferential, Lockean, 'I know not what' intention. That is, on most typical uses, terms like 'tiger' and 'water' are used to refer to some mind- and language-independent kind of thing or stuff, the precise criteria of identity for which is typically unbeknownst to speakers who nonetheless count as competent with the term.¹¹ In this case, *a priori* knowledge would be a different matter entirely. If a speaker's intention in uttering 'tiger' or 'water' is this natural kind, deferential, whatever-it-is-exactly-that-constitutes-the-real-essence-of-this, then the speaker does not have the same kind of transparent access to the content expressed by their statements, as they do in the case of social-conventional kinds. So, provided that the term in question is used as a natural kind term, then Kripke's (1972) essence-identifying, *a posteriori* necessities can occur (e.g., 'Gold is the element with atomic number 79').¹²

So, when it comes to our thought and talk about natural reality, we can discover surprising necessities, as we learn about essences and laws of nature.

Some frameworks do target objective external phenomena, as opposed to conventional categories. *A priori* knowledge here is a different beast. This is discoverer's knowledge, as distinct from maker's knowledge. Such discoveries (about gold, heat, water, and so on) are certainly not *a priori* or analytic at first, though they can get sedimented into the conceptual fabric of a framework over time. After a vague threshold has been decisively passed, it may become knowable *a priori* that whales are mammals, or that water is H₂O.

In contrast, when it comes to our thought and talk about social or conventional reality, here we have *a priori* access to the data, which is hardly mysterious or supernatural. Certain of our beliefs about all Januarys, suicides, grandmothers, widows, or bachelors are immune to counterexample, and our justification in such cases has nothing to do with empirical evidence. (Many would want to call such cases non-factual—though, again, compare the discussion of some importantly different senses of the vague term 'factual' in §4.3.)

Note that I am not suggesting or insisting that natural reality and social reality are two discrete exclusive monoliths. Some issues surely might fall into both or neither; and non-trivially distinct sub-varieties may be distinguished, within each. (For example, lots of research in psychology might be seen as focused on shades of grey between these extremes.) Just rather that, since there is more than one relevantly different sense of 'reality,' when it comes to the relations between our thought and talk and its truth-conditions, this 'no statement is true but reality makes it so' slogan needs to be handled with caution.

[§]

There are distinct sorts of frameworks, from chemistry to interior decorating, from politics to logic.¹³ Social-conventional frameworks do not have an external objective mind-independent target. The relevant categories are constituted by our thought and talk. But natural frameworks are another matter, attempting to target mind- and language-independent desiderata. In the case of natural frameworks, in general, the [EG] between meaning and extension is open, with all that entails. The meaning which any individual or community semantically associates with 'aluminum' or 'water' may not suffice to determine any specific extension, but it is a different matter entirely to press that kind of case for 'grandmother' or 'fortnight.'

Compare the sense in which the slogan 'no statement is true but reality makes it so' applies to each of the following, to get a sense of the fair degree of diversity:

1. Aluminum is a metal.
2. Cats are animals.

3. Pencils are artifacts.
4. Widows are formerly married women whose spouse has died.
5. A fortnight is a period of 14 days.

Here we see a more or less gradual transition from ‘targeting and making a judgement about’ cases, toward ‘categorizing something which is presupposed’ sort of cases. Quine’s dictum might apply to all five, to be sure, though it surely seems to do so in quite different ways to [1] and to [5].

We might also compare the above along a dimension explored by Burge (1979)—that is, openness to correction. (This will be deeply relevant to discussions about the limits of externalism to come in chapter 6.) If scientific consensus were to come to hold that [1] is false, that would be a surprise; but surely we would all fall in line, without much in the way of a shock to the rest of our world-views. (I mean, what do I know, really, about the nature of aluminum or the precise criteria for being classified as a metal?) [4] or [5] simply could not turn out to be false, as I currently use the relevant terms—though such words and concepts are prone to evolve over time. As for something like this:

6. If P and Q, then P.

In this case I have no grasp whatever on the kind of conceptual revolution it would take for me to reject this belief. Hence, these three are ranked in increasing order of resistance to correction. [1] is high on the deferential index but low on transparency; and the moves to [4] and then [6] are moves toward less deference and more transparency.

There is much more on this point, and its consequences, throughout chapter 6. We will come back to this issue and consider its connection to the Metasemantic, Semantic, and Pragmatic theses about the [EG]. The main present point is to introduce a discussion of the different ways in which Quine’s dictum applies here, across a variety of cases. The externalist challenge should be understood to apply differently to different aspects of the lexicon. There are relevant divisions which must be respected between kinds of terms, when it comes to incorporating the effects of the challenge of revisability, in addition the externalist challenge.

[§]

OK, fine, no statement is true but reality makes it so; but distinguishing these crucially different senses of ‘reality’ is essential. Both categories of natural and social reality will admit of interestingly different sub-varieties. Rather than heading down that path right now, though, the pressing ongoing task of Part III is to apply this distinction to the question of the proper

morals of the challenges of externalism and revisability, for our ongoing inquiry. Indeed, when it comes to both of these challenges, different morals and refinements are going to be appropriate, across this natural/social reality divide.

This distinction between natural and social reality, in conjunction with the work in §5.2 on frameworks and the constitutive *a priori*, highlights precisely how and why semantic and epistemic immunity to counterexample survive the challenge of revisability. It is the ever-re-evaluated frameworks that we co-construct, and to which we have non-mysterious, non-empirical access. Constitutive *a priori* knowledge and analytic truth involving the fabric of such frameworks seems to be both unproblematic and significant. Here, semantic intuition penetrates the veil (and U can reach J), whereas, in the case of natural reality, semantic intuition is subject to potential [EG] blockage (between U and J). We try to fashion our frameworks to track or reflect the nature of things as best we can, and while in retrospect we can see clear progress over time, at any point we cannot be conclusively certain that we have got things right, etc.; in contrast, social reality and our frameworks co-constitute each other. Hence, constitutive *a priori* knowledge and analytic truth involving social reality seems to be eminently attainable.

Most generally, then, the main aim of Part III is to show how an understanding-based, constitutive *a priori* orientation is well-equipped to absorb the challenges of revisability and externalism, and to afford an adequate, non-obscure answer to Plato's problem. The more specific business of this chapter has been to lay out those challenges, and to begin to develop, first, the modal revisionist response to the former and, second, the moderate externalist response to the latter. The [UJ] connections form the spine of this orientation, and the [UJ] principles capture their fundamental relevance to Plato's problem.

NOTES

1. In addition to Quine, other leading figures in work pertaining to the challenge of revisability include Mates (1950), White (1950), and Goodman (1952). A more general epistemic fallibilism of course long predates this period, dating back to ancient skepticism. More proximate influences for this above work can be found in Pierce and Dewey.

2. Kuhn (1964) is a classic early discussion; Friedman (2001) and Stump (2015) are more recent treatments which are more explicitly along the rails of a constitutive *a priori* orientation.

3. Cf. p.iv of the Preface for an overview of constitutive *a priori* views; and §6.1 for an in-depth look at three instances. Field (2000, 2005) provides an example of

modal revisionism which is not a constitutive *a priori* view (at least not explicitly, though it is consistent with one, as we will discuss in §6.3).

4. The influence here of Wittenstein (1921) is palpable. Consider 5.473 ‘In a certain sense, we cannot make mistakes in logic,’ 5.4731: ‘What makes logic *a priori* is the impossibility of illogical thought’. Certain axioms are conditions for the possibility of intelligible discourse; to change them is to change the framework of discourse itself. (Alternatively, compare what it would be like to give up on the following two beliefs—‘Neptune has four moons’ versus ‘Squares have four sides’. The former would be easy and relatively inconsequential, but the latter would involve a change of framework. The meaning of ‘Neptune’ would preserve unscathed, but not so for ‘square’!)

5. Even Maddy (2000: 109–10) throws in a couple of nice anti-Quinean, pro-revisionist cases, despite her self-styling as a Quinean naturalist. More on this in §6.3.

6. This has to be qualified in order to apply to Strawson (1959), and to any other traditional view which attempts to accommodate reference-borrowing (i.e., in which a speaker intends to exploit what others know instead of presuming an autonomous connection to the referent). However, Strawson still belongs within the traditional camp, since on his view reference-borrowing just passes the buck to some other agent. That is, Strawson’s view is that the traditional constraints need not apply to every single utterance; whereas (as we will see) Kripke’s view is that the traditional constraints are deeply misguided. (Cf. Kripke [1972: 90–92] for discussion of Strawson’s view, and Kripke [1986] for related discussion.)

7. Kripke (1972) is the most thorough and influential source here. Other important contributions include Donnellan (1970), Putnam (1975), Burge (1979), and Kaplan (1989). For recent overviews of semantic externalism, cf. Gertler (2012), Kallestrup (2012).

8. I am tempted to say this of Burge too, based on his criticisms of certain kinds of move which (as we will see below) are key to the Metasemantic thesis—cf. e.g., (2007: especially pp. 157, 160–61). However, Burge is primarily concerned with the mental, not the linguistic. So while lots of Burge (1979, 1986, 2007) obviously has deep relevance to some of these issues—to cite just one other example, Burge’s (1979) work on the significance of how the agent would respond to correction is crucial for my case in favor of the Pragmatic thesis—still I am reluctant to try to tie his anti-individualism about psychological content to any particular position on the spectrum which I draw herein, pertaining to the range of the [EG].

9. Philosophers are of course not known for their lack of audacity—cf. Williamson (2006, 2008—especially 2008: 95–96) for an attempt in the case of ‘and’, and Boghossin (2011), Sullivan (2015b), for responses. This issue is discussed in depth in §6.2.

10. The [EG] and the reference determiner are interdefinable—there is an [EG] iff there is need to posit a reference determiner, as distinct from both meaning and extension. So, both may be seen as labels for what is involved when transparency fails. (That the reference determiner is distinct from the meaning is that Frege’s [1892a] telescope is occluded, perhaps better characterized as a kaleidoscope rather than a telescope. What we have transparent access to may not limn or shadow anything beyond the bounds of our concepts.)

11. There are of course lots of non-typical uses of these terms (as of any other). Consider, for example, Chomsky's (1993, among other places) pessimism about the externalist tenet that H_2O is the essence of water, on the grounds that what counts as water in lots of places is a lot less purely H_2O than Sprite or tea is. Such uses of 'water' are not uses as deferential, essence-targeting natural kind terms, but rather uses as practical kind terms (i.e., 'whatever it is that flows out of this tap'). Cf. below §6.4, especially note 18.

12. See Sullivan (2003b, 2012) for discussion of the notion of deference at work here. It will also be further developed below, especially in §6.4.

13. There are parallels between this fissure between natural versus conventional frameworks and the ever-present realism versus constructivism debates. Realism is to natural reality what constructivism is to conventional reality. Realists are at their most realist when it comes to unspoiled, pre-categorized nature; whereas constructivists take the cognitive activity of categorization to co-construct and constitute reality.

Chapter 6

Modal Revisionism and Moderate Externalism

The ultimate goal of Part III is to further ground and develop the constitutive *a priori* as a comprehensive, satisfactory stance on Plato's problem. chapter 5 has gotten the ball rolling, when it comes to how to absorb revisability and externalism into a constitutive *a priori* framework. The aim of this chapter is to further develop the contours of my modal revisionist response to the challenge of revisability and of my moderate externalist response to the externalist' challenge. Further, the significant confluences between these responses will be highlighted and drawn out.

I begin in §6.1 with a more detailed sketch of some (connected, overlapping) versions of the constitutive *a priori* orientation. Next, in §6.2, I give a thorough treatment of a serious challenge to that orientation. Then §6.3 integrates the revisionist response to the challenge of revisability, §6.4 covers some terrain which overlaps both revisability and externalism, and §6.5 develops the moderate response to the externalist challenge. Finally, §6.6 concludes Part III by mapping my own favored version of the constitutive *a priori* view.

§6.1: SOME VERSIONS OF THE CONSTITUTIVE *A PRIORI*

I will distinguish three (overlapping, related) versions of the constitutive *a priori* orientation as (i) the Wittgensteinian, (ii) the Carnapian, and (iii) the Kantian. (That is a bit awkward, as (i) and (ii) contain heavy doses of Kantianism; but the third is distinguished by its more explicit commitment to transcendental philosophizing.) Among the core ties that bind them all together are, first, the idea that a *a priori* is at least in part a matter of status,

function, or role, as opposed to marking off some categorically different kind of content or faculty of mind, and, second, that it is a necessary condition for any inquiry that some rules or principles have some such status.

[§]

Our story begins with Wittgenstein's struggles within the philosophy of logic, prior to the World War I, and traces the continued development of the core idea that there is a crucial difference in status between 'rules' and 'propositions.' The latter bookend will be provided by excerpts from Wittgenstein (1969), his most sustained treatment of epistemological questions, written in the last 18 months of his life. The present aim is to document Wittgenstein's seminal and enduring commitment, stretching between those bookends, to a constitutive *a priori* approach to Plato's problem.¹

It may strike a discordant note to portray this most anti-historical and non-traditional of philosophers as struggling with Plato's problem. However, Wittgenstein certainly engaged—throughout his entire career—with the notion of immunity to counterexample as it crops up in the philosophy of logic. In the *Notebooks 1914–1916* (1/6/15), Wittgenstein puts this struggle in familiar terms: 'The great problem around which everything that I write turns is: Is there an order in the world *a priori*, and if so what does it consist in?' Like Kant (even if not explicitly inspired by Kant), he develops an original answer to that question, which self-consciously differs from the more familiar varieties of rationalism, empiricism, and skepticism.

I quote next an excerpt from Coffa (1991: 163), which describes some themes which are germane to this ongoing discussion:

What justifies the logical laws and rules? ... In Wittgenstein's view, an accurate diagnosis of the situation has been prevented by the propositionalist prejudice, by the insistence on looking at logic and the other *a priori* disciplines as being expressed in statements that convey facts, just like any *a posteriori* statement—except that the facts in question are somehow otherworldly. If logical laws were essentially like all other statements, they would require some sort of justification, as every statement does. ... Wittgenstein's solution is to say that logic ... has no justification [in the sense that] it cannot be conveyed by means of claims. There isn't a 'hard' domain of *a priori* truth and a 'soft' domain of *a posteriori* truth. ... This does not eliminate the *a priori* but rather locates it as a 'hardness in the soft.'

The allusion toward the end is to another quote from the *Notebooks* (1/5/15): 'My method is not to sunder the hard from the soft, but to see the hardness in the soft.' Plato's dualistic metaphysics is the classic example of sundering the hard from the soft, and Russell (1912) supplies some fairly pristine examples of its continuing prevalence (one instance of which was cited in §4.2).

Let's take 'the propositionalist prejudice', as it pertains to the philosophy of logic, to be the view that all claims are to be treated equally, when it comes to justification—that is, they are justified iff they correspond to some mind- and language-independent fact, which is their linguistically determined truth-condition. With a few pinches of salt, propositionalist justification can work for many varieties of statement ('The cat is on the mat,' '2 + 2 = 4,' 'F = MA,' 'Murder is wrong'). However, in the philosophy of logic, it forces a (familiar enough) choice between an obscure Platonism and an unsatisfactory relativism. For cases like &-elimination ($\Phi \ \& \ \Psi, \ \therefore \Phi$), for example, it seems evident that we are justified in believing it to be completely immune to counterexample. Is that because of some eternal Platonic super-fact? Or is that semblance mere illusion?

Wittgenstein charts a third option. What if &-elimination is a rule, not a proposition? The reason why that pattern of inference will never lead you astray is not that there is anything magical about the proposition expressed, but rather that it encapsulates a rule which constitutes what it means to employ '&.' (Justification in this case is more like exceptionless generalizations regarding bishops moving diagonally in chess, than like staking a claim which may or may not turn out to be subject to counterexample.) Immunity to counterexample in logic is more a matter of consequences of implicit definitions, than the discovery of a mind- and language-independent law.² Consider a few other quotes from the *Notebooks*, articulating this idea:

(2/9/14): It must in a certain sense be impossible for us to go wrong in logic.

(8/9/14): The 'self-evidence' of which Russell has talked so much can only be dispensed with in logic if language itself prevents every logical mistake.

On into the *Tractatus* we get the following ideas, already discussed in §5.2 above:

5.473: In a certain sense, we cannot make mistakes in logic.

5.4731: What makes logic *a priori* is the impossibility of illogical thought.

Logical truth has the status of the cement holding together the foundations of a language game, of thereby defining what counts as an intelligible move within the game.

This early Wittgensteinian philosophy of logic is a clear and seminal case of the constitutive *a priori*. What is distinctive about logic is its status. Immunity to counterexample in logic is a matter of the role or function which certain axioms play in the framework; and this orientation stiff-arms many traditional metaphysical questions about the (otherworldly) nature of (certain kinds of) truth-makers. (If there is a truth-maker, it is not a hunk of being [i.e., some or

other sort of fact, or state of affairs] but rather resides in the rule's *playing a certain role* or *having a certain status*.)

Further, this thread runs throughout Wittgenstein's entire career, as the *a priori* gets embodied in the guise of 'grammar' in the middle works, and on into its most extensive development in *On Certainty*. One of the things for which that work is known is Wittgenstein's 'Moorean turn'—Wittgenstein reportedly (1969: vi) loved Moore's (1925) 'Defense of Common Sense,' and engages with that essay in various ways at various points. In terms of the *a priori*, the Moorean turn is amenable to Wittgenstein for its seeing the hardness within the soft—that is, its exploration of various senses in which various sorts of seemingly empirical matters of fact should also be seen as immune to counterexample. One of the directions in which Wittgenstein takes this line is toward the conclusion that *a priori* and *a posteriori* are not discrete monoliths, but rather come blended together in various degrees, among language games and world pictures:

- §52: There isn't a sharp boundary line between [rules and propositions].
- §308: Not everything that has the form of an empirical proposition is one.
- §309: Rule and empirical proposition merge into one another.
- §318: There is no sharp boundary between methodological propositions and propositions within a method.
- §319: But wouldn't one have to say, then, that there is no sharp boundary between propositions of logic and empirical propositions? The lack of sharpness *is* that of the boundary between *rule* and empirical proposition.
- §454: There are cases where doubt is unreasonable, but others where it seems logically impossible. And there seems to be no clear boundary between them.

On Certainty also instances, related to this Moorean turn, the second aspect of the contingent *a priori* view, in addition to this first point about status—namely, the idea that holding something to that status is a necessary condition for any inquiry. (This second aspect is especially stressed in version (iii) below, which I am calling the 'Kantian' constitutive *a priori*). Namely: it is a necessary condition for any intelligible, systematic inquiry that some things get afforded this special status, function, or role:

- §337: One cannot make experiments if there are not some things that one does not doubt. But that does not mean that one takes certain presuppositions on trust.
- §341: That is to say, the *questions* that we raise and our *doubts* depend on the fact that some propositions are exempt from doubt, and are as it were like hinges on which those turn.

§342: That is to say, it belongs to the logic of our scientific investigations that certain things are *in deed* not doubted.

§343: We just *can't* investigate everything, and for that reason we are forced to rest content with assumption. If I want the door to turn, the hinges must stay put.

§415: Certain propositions underlie all questions and all thinking.

§450: A doubt that doubted everything would not be a doubt.

§509: I really want to say that a language game is only possible if one trusts something (I did not say 'can trust something').

This is a core pillar of the constitutive *a priori* orientation—that is, that treating some things as having the status of immunity to empirical counter-example is a necessary condition for any rational inquiry. This is perhaps most stark (though as always, under-explained) here:

§425: Although 'I cannot be making a *mistake* about it' that still does not entail that 'I am infallible about it'.

(This is similar to what Poincare says about the status of a definition in geometry, as we will again see below.)

Relatedly, the [UJ] connections, which form the spine of understanding-based approaches to *a priority*, are also evident in *On Certainty*:

§80: The *truth* of my statements is the test of my *understanding* of these statements.

§83: The *truth* of certain empirical propositions belongs to our frame of reference.

§§96–99 blends all of these diverse strands together, in the course of developing some of the most widely cited metaphors from *On Certainty*:

It might be imagined that some propositions, of the form of empirical propositions, were hardened and functioned as channels for such empirical propositions as were not hardened but fluid; and that this relation altered with time, in that the fluid propositions hardened, and hard ones became fluid. The ... river-bed of thoughts may shift. But I distinguish between the movement of the waters on the river-bed and the shift of the bed itself; though there is not a sharp division of the one from the other.

But if someone were to say 'So logic too is an empirical science' he would be wrong. Yet this is right: the same proposition may be treated at one time as something to test by experience, at another as a rule of testing. And the bank of the river consists partly of hard rock, subject to no alteration or only to an

imperceptible one, partly of sand, which now in one place now in another get washed away, or deposited.

To summarize: one thread which runs through the entirety of Wittgenstein's career is a seminal commitment to a constitutive *a priori* orientation, when it comes to justification in the philosophy of logic. Wittgenstein is clearly a proponent of both, first, the idea that *a priori* is (at least in part) a matter of status, function, or role, as opposed to marking off some categorically different kind of content or faculty of mind, and, second, that it is a necessary condition for any inquiry that some rules or principles have some such status. There is a strong dose of context-sensitive framework-relativity evident in Wittgenstein's distinctive orientation towards Plato's problem.

One yawning avenue into which I have not yet turned is, namely, that the specter of relativism haunts—that is, the worry that Wittgenstein's later philosophy is inhospitable to, or even incompatible with, the very idea of mind- and language-independent objective truth. I should not wade into that corner of Wittgenstein scholarship here—it is perhaps enough to say that while there is ample textual evidence of a variety of relativism in Wittgenstein (1969, as in some other places), there are also fairly categorical disavowals of any strong form of relativism (e.g., 1969: §§108, 317, 336).³ In any case, the specter of relativism haunts all non-absolutists. Our primary interest in this visit with Wittgenstein is his relevance to the constitutive *a priori* position on the questions of metaphysical, semantic, and epistemic immunity to counterexample. Elements of Wittgenstein's life-long engagement with the philosophy of logic will be seen to permeate and influence lots of subsequent work on those matters.

[§]

The two other versions of the constitutive *a priori* picture are also closely related in various ways, but are rooted more firmly in the philosophy of science. (Wittgenstein's anti-scientism was notorious, and another constant throughout his complex career.⁴) In particular, the drastic episodes in the history of ideas known as scientific revolutions provide some clear examples of the more mundane phenomenon which I am calling conceptual evolution.⁵ Reichenbach (1920) and Pap (1946) are two early proponents of a certain neo-Kantian line of response to scientific revolutions which bear the key hallmarks of the constitutive *a priori* orientation. This line develops into Carnap's (1937, 1950) well-known account of frameworks and its attendant distinction between internal and external questions.

This terrain was already charted, in a preliminary way, above in §5.2; it will be more extensively developed throughout the rest of this chapter. The primary present point is to document the allegiance to this variety of constitutive

a priori orientation on the part of Reichenbach (1920) and Pap (1946), as well as fellow-travelers Poincare and Carnap.

Chronologically, this strand begins with Poincare's (1899, 1900) conventionalism (which, as remarked above, overlaps with Wittgenstein's philosophy of logic). Famously, Poincare rejects the idea that geometric axioms express propositions; they are rather on his view disguised definitions which are meaning-constituting. In this respect, Poincare was a proto-framework-relative constitutivist. Coffa (1991: 140) and Friedman (2000: 376; 2007: 99), for example, both cast Poincare's view as very much akin to Carnap's (1950)—that is, what appears to be an arbitrary convention from the outside at the same time appears unshakably necessary from the inside:

Thus, in the case of *a priori* claims, one and the same linguistic form may be seen as playing two radically different ... roles: When regarded from outside a linguistic framework, it must be seen as ... a definition in disguise; when regarded from within the defined framework, that very sentence now expresses a claim, one true in virtue of the constituted meanings and therefore necessary. (Coffa 1991: 140)

Like more or less any understanding-based approach to *a priority*, the constitutive *a priori* orientation certainly contains an element of this kind of conventionalism—that is, Poincare's 'convention,' akin to Wittgenstein's 'rule,' is precisely a special status or role.

Our next step, after Poincare, is Reichenbach's (1920) effort to absorb the theory of relativity into a broadly Kantian epistemology. Consider the following morals:

The doctrine of the *a priori* has been transformed into the theory that the logical construction of knowledge is determined by a special class of principles, and that this logical function singles out this class, the significance of which has nothing to do with the manner of its discovery or the duration of its validity. (p. 94)

'*A priori*' means 'before knowledge,' but not 'for all time' and not 'independent of experience'. (p. 113)

Here we see an even more explicit statement than in Poincare of the move to treating *a priority* as a matter of status, not just of content, or of faculty of mind.

Moving on into Pap (1946), who explicitly builds on both Poincare and Reichenbach, we get a kind of Wittgen-Moorian context-relativity, as the soft hardening over time:

If ... our point of view is dynamic or developmental, we shall find that what were experimental laws at one stage come to function, in virtue of extensive confirmation by experience, as analytical rules or 'conventions'. (vii)

A priori is characterized in terms of functions which propositions may perform. ... A proposition which is *a priori* in one context of inquiry may be *a posteriori* in another context. (viii)

A priori is a matter status, and status is a highly context-sensitive notion. (The importance of this temporal dimension is further explored below at §6.3, and will significantly influence the final maps to be drawn in Part IV.)

Pap segues smoothly into Carnap's (1937, 1950) better known work on framework-relativity. The conventional-from-the-outside dimension, with its attendant specter of relativism, is there loud and clear: 'In logic there are no morals' (Carnap 1937: 52), and the scientific spirit in philosophy demands tolerance as to the proliferation of linguistic frameworks (1950: 40). As Friedman (2000: 371) puts it, for Carnap at this stage: 'All standards of "correctness," "validity," and "truth" ... are relative to ... linguistic framework. ... Such rules are constitutive of the concepts of "validity" and "correctness."'

The work of explaining how and why this conventional-from-the-outside is also, seen from the inside, a robust and significant kind of immunity to counterexample a primary aim of the rest of this chapter. How and why is this not just a nihilistic relativism? How does it not fall prey to some of the daunting challenges to a pure conventionalism in the philosophy of logic (e.g., Quine [1936], Prior [1960])? Answering that will necessitate delving into the complexities of conceptual evolution, and into the shifting sands around the border between change of language and change of theory.

[§]

Version three is the most explicitly Kantian take on the idea; its most widely read proponent is Friedman (1992, 2000, 2001, 2011), but other proponents include DiPierris (1992), DiSalle (2002), Richardson (2002), and Franco (2011). While elements of Wittgenstein, Reichenbach and others discussed above are evident here as well, the explicit and self-conscious Kantian talk of transcendental philosophy as a required *a priori* supplement to scientific theories provides something of a third variation on this theme, in addition to Carnap's (1950) frameworks and Wittgenstein's (1953) language games or (1969) world pictures.

The guiding analogy is to space and time in Kant (1781), enjoying a status undreamt of by previous rationalists and empiricists—that is, transcendental preconditions for experience. Space and time, for Kant, are more like rules of the game than like Platonic otherworldly superfacts or Humean inductions from experience. To experience the happening of an event *in space and time* is, quite simply, constitutive of experiencing it at all. In the constitutive *a priori* tradition, this distinctively Kantian status is then transposed from preconditions for experience to preconditions for intelligible, tractable inquiry. Here

we see the key respect in which the constitutive *a priori* is a neo-Kantian orientation—Kant’s innovative move in epistemology is precisely about a new status, previously undreamed of in the long history of rationalism-empiricism debates.

In Friedman’s (2000, 2001) version of this general program, he distinguishes in status three different levels or components of scientific inquiry: (1) at the base level, we have concepts and principles of empirical natural science; (2) at the second level, there are the constitutively *a priori* principles, within which alone empirical testing at the base level is possible; (3) at the third level, there are philosophical meta-paradigms or meta-frameworks.⁶

The relations between (1) and (2) are familiar enough by now—as this point has been explicitly with us since at least §5.2. Other examples of the myriad ways in which some things must be held fast in order to generate and intelligibly isolate specific questions are discussed by Friedman (2000), who develops at some length the examples of the relations between the mathematics and the physics in both Newton’s and Einstein’s epochal scientific innovations. It is the (2)–(3) relations which are distinctive here, and which proponents take to be where this third Kantian sort of constitutivist moves beyond conventionalism:

It is precisely here that we need to move beyond both Carnap and Kuhn, by describing a fundamentally new function for what Kant called transcendental philosophy. (379)

The enterprise that Kant called transcendental philosophy—the project of investigating and philosophically contextualizing the most basic constitutive principles defining the fundamental framework of empirical natural science—plays an indispensable orienting role with respect to conceptual revolutions within the sciences ... thereby makes available prospective notions of inter-framework rationality. (382)

Friedman sees his Kantian constitutivism as an advance over previous versions, on which it is a further development, precisely for its stance on progressive relations between frameworks.

Relatedly, and significantly, these most recent, most Kantian constitutivists understand a large part of their job to be to explain why a seamless Quinean holism cannot afford an adequate epistemology in general, or account of the history of science in particular. Friedman (2000: 376) says: ‘It will not do ... to view ... the constitutively *a priori* parts of our scientific theories as simply relatively fixed or entrenched ... relatively difficult to revise.’ In Stump’s (2011: 188) terms: ‘The constitutive principles are not merely more entrenched—they have *never* been directly tested.’

Akin to their Wittgensteinian and Carnapian allies, then, these Kantian constitutivists in the philosophy of science also clearly adhere to these core

tenets of: (i) *a priori* is a matter of status, role, or standing, not merely a special kind of content or faculty of mind, and (ii) *a priori* thus understood is a necessary condition for any intelligible inquiry.

§6.2: A CASE STUDY IN UNDERSTANDING AND JUSTIFICATION

It is clear from chapter 5 that the [UJ] connections and principles are absolutely essential, the spine of any understanding-based approach to *a priori*. The next issue then is a long focused case study into what is perhaps the most significant prima facie counterexample, threatening to undermine these connections. Working through this challenge is not only important in itself, but will also help to hone some distantly related contours of the constitutive *a priori* orientation.

I will use Conjunction Elimination—the inference from a conjunction to one of its conjuncts—as my stock example of a pattern of inference which is safely known to be valid:

$$\begin{array}{l} [\&E] \Phi \&\Psi, \\ \therefore \Phi \end{array}$$

Hence, any specific instance of [∧E] counts as a safely known logical truth. Some of the perennial questions within the epistemology of logic are due to the sense that such knowledge exhibits a remarkable *immunity to counterexample*: that is, it is not just that I have yet to encounter a situation in which a conjunction failed to entail one of its conjuncts (which would be remarkable enough, to be sure), but, further, there is the atavistic intuition that such a scenario would be both epistemically inconceivable and metaphysically impossible. And hence, questions about the epistemology of logic are entangled with some rather large philosophical issues, such as *a priori* knowledge and necessary truth.

The [UJ] connection is perhaps most strongly evident in the case of logical truths. (Surely the claim that one who understands ‘∧’ is *thereby* justified in believing an instance of [∧E] to be a logical truth is safer than Anselm’s claims about what is entailed by understanding the concept of God!) According to this approach to the epistemology of logic, one’s justification for holding that [∧E] is a valid pattern of inference is grounded in one’s grasp of what ‘∧’ means. (Alternatively, holding that [∧E] is immune to counterexample is a necessary condition for competence with the concept of conjunction.) Versions of this understanding-based epistemology of logic are rather ubiquitous—for example, variants can be found in the work of both

Leibniz and Hume—well prior to Poincare or Wittgenstein—and instances of it are recently developed as a version of ‘rationalism’ by Peacocke (2000, 2004) and as a version of ‘empiricism’ by Boghossian (1997, 2000, 2011).

Given both the prevalence of this approach to the justification of logical truth, and the sense that logical truth is perhaps the most viable case of this core [UJ] connection, the deviant logician objection threatens to wreak considerable havoc in the house of philosophy. For the claim pressed by the proponents of the DLO is precisely that even here in the pristine confines of pure logic, understanding falls decidedly short of affording justification. (And see Williamson [2008] for a sustained attempt to draw out sweeping, revisionary meta-philosophical conclusions from the DLO.)

As preliminary, I should sketch (at least a little bit) what makes for ‘deviance’ in logic. There is general consensus as to what constitutes ‘standard’ or ‘classical’ logic, fundamental tenets cementing the foundation of the enterprise, which hold constant from Aristotelian categorical logic through (and beyond) modern propositional and predicate logic. Core here are the Law of Excluded Middle and the Law of Non-Contradiction:

$$\text{LEM} : \Phi \vee \sim \Phi$$

$$\text{LNC} : \sim (\Phi \& \sim \Phi)$$

Deviant logics are those which transgress such standard, classical tenets. Thus understood, deviant logics are hardly a novel phenomenon: the idea that LEM is subject to counterexample (for future contingents, say, or conditionals with a false antecedent) was fairly prevalent throughout Ancient and Medieval philosophy. However, the monolithic status of standard, classical logic is more drastically under siege in the current era than at any previous time. Intuitionist logics, many-valued logics, and fuzzy logics are some fairly well-known, fairly recently well-developed logics which categorically reject LEM, as do many contemporary theories of vagueness. There are paraconsistent logics which develop the idea that rejecting LNC is the best way to handle the semantic paradoxes (e.g., ‘This sentence is false’), among other phenomena, and quantum logics also reject LNC. In these liberal times, it is even fairly common and plausible to work with different logics for different purposes, in different contexts.⁷

Of course, in logic (as in life) deviance has its price. Perhaps most notably, one cannot have proof by contradiction (a.k.a. indirect derivation, *reductio ad absurdum*) without LEM, and many pillars of both logic and mathematics have as yet only been proved in this way. Before getting back to our main themes, I will quote Quine’s statement of the (open-minded but just barely so) orthodox party line on this question:

[L]et us not underestimate the price of deviant logic. There is a serious loss of simplicity, especially when the new logic is not even ... truth-functional. ... [T]he price is perhaps not prohibitive, but the returns had better be good. (1970: 86)

[§]

The primary target of the DLO, then, is this general category of views which ground justification for logical truths in understanding their components:

[UJLT]: our justification for logical truths is grounded in our understanding of their constituent concepts

(This is a relatively bulletproof, bedrock instance of the [UJ]s.) The strategy is to undermine [UJLT] by offering counterexamples to the following putative corollary:

[UJLT corollary]: competent agents who share the same understandings of logical primitives could not coherently disagree as to whether something constructed out of commonly-shared primitives counts as a logical truth.⁸

To the extent that one can motivate the notion of disagreement about what ought to be counted as logical truths, among those with a shared understanding of the meanings of the constituent primitives, one thereby motivates skepticism that understanding could suffice for justification.

Now, one possible line of response to the DLO would be to question whether [UJLT corollary] really is entailed by [UJLT]. However, for present purposes, I will concede the corollary. *Prima facie*, it seems that proponents of [UJLT] must classify any disagreement as to whether something ought to count as a logical truth as ultimately stemming from one of the following two sources: (i) at least one party falls short of a competent, comprehensive grasp of one of the relevant concepts, or (ii) at least one of the relevant concepts is ambiguous, understood in different senses by the different parties. In referring back to these, I will call (i) ‘the incompetence option,’ and (ii) ‘the ambiguity option.’

All versions of the DLO, then, argue that there can be disagreements about logical truth which involve neither incompetence nor equivocation. One version of the objection can be found in Horwich (2000: 158–59, 2006: Ch. 6), focused on disputes between an intuitionist and a classical logician about whether instances of LEM should be counted as logical truths. I take it that it would not be remotely satisfying for a defender of [UJLT] to avail of the incompetence option—that is, to just insist that the intuitionist *ipso facto* lacks a competent grasp of the classical concepts of negation or disjunction.

To the contrary, the intuitionist challenge presupposes a grasp of those concepts, and rejects some of their relatively unpalatable consequences.

A defender of [UJLT] might attempt to take refuge in the ambiguity option—that is, the idea that, as a result of their disagreements about what ought to count as valid, the intuitionist ends up with distinct concepts of negation, disjunction, etc. However, this option too encounters some complications. For example, the ambiguity objection does not seem to be strong enough to quell the worry. What resources does the classical logician have to handle an intuitionist who obstinately insists, in the face of the ambiguity objection, ‘NO! I mean exactly what you do by the terms “negation” and “disjunction”?’ Nothing, it seems, but the fallback to the incompetence option, which we already found to be wanting.⁹

So, what does Horwich’s deviant logician show? Does reflection on the intuitionist challenge to classical logic show up something deeply suspicious about [UJLT], and, more generally, about the alleged core [UJ] connection?

Versions of the DLO are developed more thoroughly by Williamson (2006: §2, 2008: Ch.4). Williamson argues that logically competent agents can even have unequivocal, informed, engaged disagreements about whether something of the form ‘All As are As’ is an instance of a logical truth. For example, there are (sophisticated, considered) reasons to worry about the relations between existential import and truth-conditions—that is, to hold that any statement that purports to refer to ‘A’s can only be true if there exist As (in the relevant context). To the extent skepticism about the existence of As can be motivated, then we can imagine someone with (sophisticated, considered) reservations about whether a particular instance of ‘All As are As’ should be counted as a logical truth. (E.g., Is ‘All unicorns are unicorns’ a logical truth?) Again, the incompetence and the ambiguity options do not have much promise to handle all possible dissenters (as Williamson argues).¹⁰

Another example developed by Williamson concerns the logic of vagueness. Many theories of vagueness posit truth-value gaps. To the extent that one can motivate the claim that ‘A’ is vague, to proponents of such a theory, then, again, we can imagine someone whose considered judgment is to balk at whether a particular instance of ‘All As are As’ (e.g., ‘All tall people are tall’) should be counted as a logical truth. Again, Williamson argues that neither the incompetence nor the ambiguity option can save [UJLT] on this front.

At the same time, Williamson holds (as do I) that any statement of the form ‘All As are As’ is a logical truth. Hence, Williamson believes he has provided counterexamples to [UJLT corollary]—that is, logically competent agents who understand instances of logical truth, but yet do not assent to them. So, then: Does the DLO show that there can be no deep constitutive connection between *understanding* and *justification*, even in the relatively straightforward case of logical truth? Does the DLO prove, a fortiori, that

nothing can have the status such that assenting to it is a necessary condition for understanding it?¹¹

[§]

What then should a proponent of the constitutive *a priori* orientation say about this challenge? Well, they hold that it is an important lesson of mid-twentieth-century modal epistemology—a legacy of Wittgenstein, Carnap, and others—that such notions as *a priority* are comprised of two distinct, separable factors. *A priority* depends on both intrinsic content and place in a framework. If we appreciate this point, then the fact that two competent agents could unequivocally agree about something’s intrinsic content, while attaching different statuses to it, is no knock-down challenge to core [UJ] connections. Horwich’s and Williamson’s challenges pose no more of a problem for [UJLT], for constitutivists, than does the fact that two agents might agree that every event has a cause, and yet for one of them this is an *a priori* regulative rule while for the other it is an *a posteriori* inductive generalization (cf. §5.2).

In other words, for proponents of constitutive *a priori* orientation, there are at least two separable factors which constitute ‘shared understanding’—that is, sameness of content, and sameness of status. Hence, this orientation offers proponents of the [UJ] connections a principled defense from the DLO. For it offers a way to articulate and develop the intuition (which is no doubt motivating the authors alluded to in note 8) that deviant logicians do not in fact instance a ‘shared understanding’ of the basic tenets of standard logic (as is demanded by any notion that the DLO undermines [UJLT]). So, even if their conceptions of the content are identical, their attaching different (and deviant) status to that content undermines their promise to afford counterexamples to [UJLT].

Consider again the intuitionist challenge to LEM. LEM is a fruitful simplifying element of many branches of logic; but there are deep and ancient reasons to countenance counter-instances. This situation is completely amenable to an understanding-based, constitutivist-style explanation. There are external (instrumental, conventional, pragmatic, etc.) questions about what we want from a logic. Again, these days it is even fairly plausible and common to work with different logics for different purposes and contexts, where the appropriate external questions vary from case to case. For some of these external questions, simplicity and fecundity will receive a high ranking, and LEM has proven to be expedient toward those ends. For other external questions, comprehensiveness and integrity may trump simplicity and fecundity, and, accordingly, the putative counter-instances to LEM may be judged to be decisive.

The important point for present purposes is that intuitionism is a distinct framework from classical logic, as befits its different answers to external

questions about logic. Hence, for proponents of the constitutive *a priori* orientation, that the LEM is a standard, classical logical truth (i.e., derivable as a theorem in frameworks which adhere to the traditional answers to external questions about logic) while being invalid within intuitionist logics is no serious challenge to the core, constitutive [UJ] connections. Common understanding demands the sharing of both status and content, not just of content. Since, accordingly, the understanding of LEM varies between intuitionists and classical logicians, it is unsurprising and inconsequential that its justifiability (or lack thereof) does too.

Given certain aims and interests, it is reasonable to take LEM as *a priori* (i.e., like the principle of sufficient reason in some frameworks, simply not subject to empirical counterexample).¹² However, when it comes to [UJ] connections, there is little to be said in favor of the claim that assenting to LEM is a necessary condition for understanding it. Contrast this with [&E], where this tight [UJ] (or meaning-constituting) connection is unshakably evident (despite Williamson's audacious effort mentioned in note 11). Espousers of [UJLT] can and should treat LEM and [&E] differently, because of this clear difference.

What about Williamson's deviant logicians? Again, proponents of the constitutive *a priori* orientation should take the moral to be that foundational questions about the relevant framework are conceptually prior to what turns out to be constitutive *a priori* within any specific framework. Williamson's deviant logicians are, like the intuitionist, working within a non-standard framework, which rejects certain canonical answers to external questions about logic. One exceedingly complex challenge stemming from the development of deviant logics is that it is difficult to judiciously and comprehensively settle external questions about logic; nonetheless, this defense of the core [UJ] connections relies only on the crucial distinction, in the epistemology of logic, between internal and external questions, not on any specific answers to the external questions.

Distinct logics of vagueness (say) will countenance disjoint sets of logical truths. It is not a trivial matter to decide which logic of vagueness one ought to prefer, all things considered; but constituent [UJ] connections are a separate matter. Once we have settled, however tentatively, what we want from a logic of vagueness, we will accordingly settle on answers to the appropriate external questions. In due course, there will issue framework-relative logical truths, such that justification for them is grounded in understanding them. One is, of course, free to deviate; but one thereby changes the framework.

Hence, proponents of the constitutive *a priori* need not take the DLO to have undermined the core [UJ] connections, because they can explain how it is that deviant logicians do not share an understanding of the relevant tenets

with their classical opponents. For example, there is a clear sense in which an intuitionist and a classical logician do not instance a ‘shared understanding’ of LEM—even if (contra note 9) they semantically associate exactly the same content with the formula ‘ $\Phi \vee \sim\Phi$.’ Likewise, while I concede that the prospects are dismal for dismissing Williamson’s deviant logicians (pertaining to ‘All As are As’) on grounds of incompetence or ambiguity, still these cases instance such different orientations with respect to external questions about logic, and attendant differences in status as to these instances of ‘All As are As,’ that constitutivists should hold that these are not instances of ‘shared understanding’—again, as the DLO needs them to be.

[§]

What deviant logicians show about the epistemology of logic, then, is that proponents of understanding-based accounts of justification for logical truths are well-advised to endorse the constitutive *a priori*. There is a distinctive notion of *a priori* (of which logical truth is a distinctive case in point) which is well-equipped to meet the deviant logician’s challenge.

The key step in any version of the DLO is to argue that there can be disagreements about logical truth between competent agents who share the same understandings of logical primitives. However, if we concede that *a priori* depends on both intrinsic content and place in a framework, then the fact that two competent agents could unequivocally agree about something’s intrinsic content, while attaching different statuses to it, is neither surprising nor disturbing. If logical truths are constitutive *a priori* truths, two competent agents could unequivocally agree about something’s intrinsic content, while attaching different statuses to it. Hence, then, the DLO does not afford counterexamples to [UJLT].

To treat something as *a priori*, as simply not possibly subject to empirical disconfirmation, is to mark off a certain content as having a certain privileged status. I have been using [&E] as a case where the intrinsic content all but guarantees the special modal status; but such meaning-constituting cases are relatively rare. The case of LEM helps to illustrate what a daunting job it is to get from [&E] to a comprehensive epistemology of logic, let alone to a comprehensive account of *a priori* justification. For now, though, the moral is that for proponents of the constitutive *a priori*, the possibility of a deep constitutive connection between *understanding* and *justification* in (and beyond) the epistemology of logic survives the challenge of the deviant logician.

What the prevalence of deviant logics shows is not that the core [UJ] connections are completely untenable, but rather that it is untenable to approach the notion of ‘logical truth’ as if it designates a monolithic block of eternal superfacts.

Hence understanding-based accounts of *a priori* in general, and constitutive *a priori* orientations in particular, have not been besmirched on this core front. The [UJ] connections still stand. Furthermore, and to the contrary, working through these details has just added to ongoing amassing of considerations in favor of this orientation. On then with the task of showing how this orientation provides responses to the challenges of revisability and externalism.

§6.3: REVISABILITY AND CONCEPTUAL EVOLUTION

So: that defense of the [UJ] connections on a constitutive *a priori* picture is a main plank in the main business of this present chapter. Again, my own constitutive *a priori* orientation is made up of a modal revisionist response to the challenge of revisability, and a moderate externalist response to the externalist challenge. This present section will continue this development and defense work, in ways focused most specifically on modal revisionism; §6.4 will segue into issues that pertain to both revisability and externalism; §6.5 will move on to issues which pertain more squarely to moderate externalism; and §6.6 will come back to the two challenges' confluence and collective upshot.

The complexities of change over time need to be taken into account, when it comes to the modal revisionist's framework-relative answer to Plato's problem, and the ways in which this consideration differs for the cases of metaphysical, semantic, and epistemic immunity to counterexample. Not only are multiple frameworks generally always applicable, when it comes to more or less any isolated question or embedded inquiry, but more or less all frameworks are more or less constantly under revision. Scientific revolutions may be rare and drastic, but conceptual evolution is more on the order of death and taxes. (Any scientific revolution would involve lots of conceptual evolution, but, for the most part, not vice versa.) This is a crucial aspect to be reckoned with and rendered, when it comes to articulating how it is that immunity to counterexample is retained, on this orientation.

Paying heed to this historical dimension brings into clearer relief some of the complexities within a satisfactory constitutive *a priori* picture. It is easy enough to distinguish paradigm cases of internal questions from external ones, or as Carnap later puts it, change of language from change of theory:

I should make a distinction between two kinds of readjustment in case of conflict with experience, namely, between a change in the language, and a mere change in, or addition of, a truth-value ascribed to an indeterminate statement. ... A change of the first kind constitutes a radical alteration, sometimes a revolution, and it occurs only at certain historically decisive points in the development

of science. On the other hand, changes of the second sort occur every minute.
(1963: 921)

However, in the course of ongoing inquiry there is clearly a continuum between these extremes. The re-categorization of whales, or the spitting of the atom, involved both change of language and change of theory, and neither can be seen as an easily discretely separable sub-component.¹³

Both theories and languages change over time, which involves, and prompts, change of meanings, and of the frameworks they compose. This is a glaring and significant upshot of the challenge of revisability, which it is crucial to sort out in this present project. Since languages are organic entities which change over time, that will have palpable effects on precisely what it is to understand the meaning of a given word (e.g., ‘whale’ or ‘atom’). In turn, any effect on what it takes to count as understanding a word or concept is going to have knock-on effects to any [UJ] connections or principles.

The limit of revisability is precisely the sense in which immunity to counterexample is retained. Changes as to what counts as analytic or *a priori* can only occur where there exists conceptual evolution, which in turn suffices for a change of framework. Within a framework, though, on the constitutivists’ view, immunity is not only always possible but, further, required. (As broached above in §6.1, Poincare engaged with this Janus-faced quality, whereby what appears to be malleable from one perspective can seem mandatory and unshakeable when approached from another vantage point.) One task of Part IV will be to chart the metaphysical, semantic, and epistemic consequences, when it comes to immunity to counterexample, of this dimension of change over time.

While conceptual evolution is most obviously and directly pertinent to the challenge of revisability, it will also have some deep bearing on the externalist challenge. Following up on §5.4, for starters, conceptual evolution is bound to be a very different thing for the case of natural versus conventional frameworks. In general, anything which depends on concepts will be affected by conceptual evolution; any categorical or principled difference between sorts of concepts (i.e., natural vs. conventional) will track or mark differences in these sorts of effects. Conceptual evolution in cases like ‘whale,’ ‘planet,’ or ‘metal’ is in myriad ways different from cases like ‘marriage’ or ‘money’ (i.e., our views about this mind-independent phenomenon are changing in light of new discoveries vs. if there is consensus to change this institution, then so be it).

[§]

I will next return to one of our ongoing themes of sorting through various senses in which something could be said to be immune to counterexample,

and further unpack some differences between metaphysical modality, on the one hand, and semantic and epistemic modality, on the other, when it comes to the challenge of revisability.

As detailed above in Parts I–II, that which is necessary is not revisable, for it is precisely the things which are the most firmly bolted down that that concept is used to single out. (Of course, any particular agent's guesses as to what is necessary may be revised over time, as a function of evidence, insight, etc.; but the *truth-conditions* of 'It is necessary that P' do not change over time.) However, for finite, fallible agents like us, it is only prudent good sense to hold that what counts as justified *a priori* will vary as do our beliefs and the concepts which they involve. So, for a post-Kripkean framework-revisionist, a major difference between attributions of metaphysical modality (e.g., 'It is necessary that P') and attributions of epistemic modality (e.g., 'My belief that P is justified *a priori*') is that the truth-value can change over time only in the epistemic case. Metaphysical necessity is framework-independent, whereas *a priori* justification is framework-relative. (Semantic modality is also framework-relative, for similar reasons; but here let us focus on the contrast between metaphysical and epistemic modality.)

For example, if it is necessary that heat is the motion of molecules, then it did not just become necessary once it occurred to someone, or once it became sufficiently verified and accepted by experts. Metaphysical necessities could not have been otherwise, and are categorically indifferent to whether anyone knows anything about them. In contrast, my favored revisionism holds that truth-conditions could change over time, for cases like 'It is justified *a priori* that whales are fish' and 'It is justified *a priori* that atoms are indivisible.' Note though that change in the truth-conditions of an attribution of epistemic modality can occur only given an updating of the relevant framework—that is, such a change depends on a significant change to the content of at least one of the relevant constituent concepts. This explains the qualified but still significant sense in which immunity to counterexample is retained on this sort of revisionism. My ancestors' belief that whales are fish has not been falsified, because whales still do satisfy the relevant superficial, unscientific criteria—that is, their concept of 'fish' was something like 'anything that lives in water, swims, and is more-or-less shaped like a tuna.' But, in my dialect, it is not true, let alone justified *a priori*, that whales are fish. My concept of 'fish' is a rather distinct ancestor of theirs, a distinct biological kind term, which differs in conditions of satisfaction.

So, Kripke's distinction between necessity and *a priori* lends credence to the revisionists' tenet that epistemic modalities are revisable, framework-relative, subject to re-evaluation in cases of conceptual change. Thus, absolutism about epistemic modalities is not just immodest; further, Carnap (1950), Quine (1951), Wittgenstein (1969), and Kripke (1972) converge in showing absolutism to betray a deep and arrogant mistake. Our beliefs get updated,

our concepts evolve; it is to be expected that any non-empirical (semantic, conceptual, rational, etc.) sources of justification are directly and thoroughly effected by such epistemic developments and conceptual evolution.

[§]

In addition to this dynamic source of complexity, which is the main present focus (i.e., frameworks are more or less constantly under revision) another source of complexity for this constitutive *a priori* picture is the purely static consideration that multiple frameworks might simultaneously apply in the same inquiry (whether we are talking about an isolated individual at a given time, or across a community of agents). Generally, out here in the real world, both the static and the dynamic sources of complexity (among other things) are always at play. The upshot is a deep and through context-sensitivity, which pervades and motivates the framework-relative approach to *a priority*. Railton (2000) is instructive on this front, going beyond Wittgenstein's own complex position on the context-sensitivity of the rule/proposition boundary. Railton's is a multi-dimensional dynamism, in which 'norms' are propositions given the status of rules for the purposes at hands, and 'supernorms' are created whenever there is systematic rational deliberation about norms. Neither norms or supernorms are ways of classifying content, per se, and neither is context-independent. (It is a question of how it is used, not of what it says or how you came to believe it.) There are multiple shades of grey and directions of fit.¹⁴

This static kind of factor is most familiar perhaps from the literature on scientific revolutions, in cases of severe Kuhnian crises. Which theoretical lens is brought to bear on the data will significantly impact how we conceive of its nature and upshot. Precisely what the question at issue *is* framework-relative, let alone how it should be answered. Consider some of the cases already discussed, to motivate the two separable dimensions of *a priority* (i.e., content vs. status), such as 'all events are caused' from §5.2 (between an agent who takes that to be a regulative rule, and an agent who takes it to be an inductive generalization), or LEM from §6.2 (between an intuitionist and a classical logician). As we saw, theories in which the notion of status plays a key role complicate what counts as 'shared understanding'; this will also drastically affect the scope and range of the core [UJ] principles.

Another familiar way in which this static complexity factor plays out in the literature is in the guise of Quine/Duhem holism (cf. Quine [1951], Quine & Ullian [1970]), which is particularly influential among naturalists. A guiding idea here is that there are many ways to iron in recalcitrant data (e.g., an unexpected and problematic experimental result), and hence that the unit of experiential confirmation not an individual statement. Down this avenue, shades of grey stretch between updating an isolated, individual belief versus

rejecting an entire theory. The web of belief is a compelling metaphor, for such cases. We're all repairing our ship while out at sea, after all—to borrow the metaphor which Quine borrows from Otto Neurath—at least those of us who are not absolutists. Here again we steer into the inevitability of jettisoning any pretense of a firm distinction between change of language and change of theory. Conceptual evolution essentially involves the former, and is also an inevitable component of the latter.

The static source of complexity flows confluent with the temporal, dynamic source. With conceptual evolution (and constitutive thereof) is alternation of framework; attendant upon that too is change with respect to what gets treated as immune to counterexample, what counts as analytic and hence *a priori*. Only in hindsight can conceptual evolution be neatly sorted—and, at that, that must be to some degree stipulative.

[§]

I close this section by honing my own constitutive *a priori* position on conceptual evolution via quick comparison with three assorted authors, from outside this tradition—Peacocke's (2000, 2004) program of moderate rationalism, Field's (2000, 2005) evaluationism, and Maddy's (2000, 2007) naturalism.

Peacocke's self-styled 'moderate rationalism' is an understanding-based account of *a priority* which aims to fashion an adequate account of the *a priori* without taking on any anti-naturalistic commitments associated with the acquaintance-based tradition. The notion of 'meaning-constitution' plays a key role in his program. Compare again:

[1] Squares have four sides.

[2] Neptune has four moons.

[1] is an example of a meaning-constituting belief. Hence, Peacocke's 'meaning-constitution' is very close to my [UJ], an understanding-based approach to immunity to counterexample. It bears many similar influences, from Poincaré through Carnap.

However, Peacocke's *a priori* is more of a special kind of content, as opposed to the constitutivists' status. This makes Peacocke more explicitly a kind of rationalist than those of us in the constitutive *a priori* camp, who privilege the importance of status, in addition to content, when it comes to *a priority*.

Next, Field's (2000, 2005) evaluationism is a novel, status-conscious, stance on *a priority*. The core idea is that, while most of our beliefs are 'non-basic' in the sense that they can be justified by appeal to other beliefs (tenets rules principles etc.), not all of our beliefs can have this status, or else the process of justification would be circular, or would go on forever. Some of our beliefs

must be basic, and to them we are entitled ‘by default,’ not because we can justify them by argument. There is nothing mysterious or obscure about default entitlement; it is an ‘attitude of approval’ that is accorded to certain core constituents of our ‘logic and methodology’ (2005: 86).

Now, this is very similar to the constitutive *a priori* orientation. It too is a modal revisionist response to the challenge of revisability; further, it also endorses the two prongs of the constitutive *a priori* view—that is, the status-orientation, and the point about necessary preconditions for inquiry. It may not be fair to force Field into the constitutivists camp, whether he likes it or not; but, given the above, it is tempting to take the arguments he gives for evaluationism as also further bolstering the case for the constitutivist status-conscious orientation on *a priority*.

Finally, Maddy (2000, 2007) styles herself as a naturalist heir of Quine. However, as I have set things up here, the particular substance of her explicit departure from Quine is quite significant (cf. Maddy [2000: 108ff, especially pp. 112–14]). Crucially, there is a seam in her web of belief, a ‘brute methodological distinction’ (113), of the exactly sort that has been familiar to us since §5.2 (e.g., rule/proposition, or status/content). Though not by that name, the constitutivist status of *a priority*—a very familiar kind of immunity to counterexample, for certain special aspects of theory—is alive and well within Maddy’s self-styled naturalism. Because of this, rather than countering the anti-naturalist arguments of the Kantian constitutivists, Maddy ends up espousing a terminological variant of the view.¹⁵

§6.4: SEMANTIC DEFERENCE

Next on the agenda is to delve into the notion of semantic deference, which plays a role in most recent conception of language, which appreciates its communal dimensions. An individual language user is a part of a system, which enhances the conceptual resources available to the individual. Consider our evident abilities to think and talk about Ancient Greek politics or the geography of Mars, or all manner of things with which no one I have ever met has had any first-hand acquaintance—let alone non-actual or impossible objects, which are rather thoroughly unacquaintable. Semantic deference is the appropriate attitude of the individual speaker toward the community-wide communicative powers of the terms in their thought and talk. (Natural kind terms provide a nice clear paradigm case, to which we turn below—for example, I believe that aluminum is a metal, would assent to that and bet on it and etc., and yet have no real capacity to divide either aluminum or metals from any close competitors. My attitude toward the criteria for the correct application of both terms ‘aluminum’ and ‘metal’ is deferential.)

Properly excavated, semantic deference promises to have some deep relevance to both revisability and externalism—let alone to many other ongoing discussions, about whales and atoms and the broad distinctions between natural and social-conventional reality.

As a point of contrast with some of the sorts of cases to be considered below, let us next delve a little further into the notion of semantic deference as it applies to ordinary speakers' usage of natural kind terms. For ordinary (i.e., non-expert) speakers, to use a term as a natural kind term involves a Lockean, 'I know not what,' deferential intention.¹⁶ That is, on such uses, terms like 'tiger' and 'water' are used to refer to a mind- and language-independent kind of thing or stuff, the precise criteria of identity for which is typically unknown to speakers who nonetheless count as competent with the term. (As came up in §5.4 [cf. especially note 11] and will be further explored below, there are of course lots of non-typical uses of these terms [as of others].) Typically, and tellingly, such non-expert speakers are relatively open to correction by experts, when it comes to the exact criteria for the correct application of such terms. Clearly, in these cases of deferential uses of natural kind terms there exists an externalists' gap [EG] between the conditions which typical individual speakers associate with the terms, and their actual precise technical conditions for correct application.

(Herein lies another way in to why natural kind terms play such an important role in recent work in the theory of reference, alongside the perennial model case of proper names. Using a term as a natural kind term is the very paradigm of semantic deference, and semantic deference is essential to semantic externalism.)

There are, however, stretches of the lexicon concerning which this [EG] does not seem to exist. To take an extreme example, consider the word 'and,' whose meaning I will take to be constituted by the standard introduction and elimination rules. The traditional motivations for semantic internalism still seem to be completely applicable here—that is, competence with 'and' is a matter of grasp of meaning, and that meaning determines the term's extension. So, here, we do not get the [EG] between what is constitutive of competence and what determines the extension. When I encounter the conjunction of two propositions P, Q into 'P and Q,' the semantic contribution of 'and' is completely transparent to me. As opposed to the case of natural kinds, I am autonomous, not the least bit deferential, about the semantics of 'and.'¹⁷

Though 'and' stands at the extreme, in this respect, it is not hard to find kinds of expression which are more like 'and' than like 'tiger.' Terms of elementary arithmetic, such as 'equals' and 'even number' provide strong candidates, as do simple geometric terms like 'square.' Even further, though, many common social kind terms, which target categorizations which humans construct, as opposed to discover, also provide plausible instances. So, for

example, the content of my ‘grandmother,’ ‘bachelor,’ or ‘widow’ thoughts and utterances is more autonomous and transparent than deferential. Here again the [EG] between what typical speakers typically associate with the term and what determines the extension does not get much purchase. That which is constitutive of competence is also that which determines the extension, and so content is transparent to competent speakers.

Tying back in to some points first mentioned in §5.4, degree of resistance to correction is a primary indicator of where a particular use of a particular term belongs on this deference-transparency spectrum. And so compare what it would take to convince you to reject the following:

1. Aluminum is a metal.
2. Widows are formerly married women whose spouse has died.
3. If P and Q, then P.

The move from a paradigm natural kind case to a paradigm social kind case is pretty drastic, along this index, as is the subsequent move from there to a meaning-constituting axiom of elementary logic. These are steps further and further into the un-black-swannable terrain, where the [UJ] connections and principles hold fast. Instead of an [EG], in these latter cases, we have a rather fast and tight link between meaning and extension. Its relevance to immunity to counterexample is glaring.

This line of thought requires careful handling and extensive refinement; I will begin this job immediately below. Still, though, it runs clearly counter to the metaseantic thesis about the [EG] (i.e., the idea that there is an [EG] for all types of usage of all types of term). Statements [2] and [3] provide examples of mundane counterexamples to such extreme or unqualified claims about competence, transparency, and externalism (which were discussed above in §5.3, and to which we will return below in §6.5). Since the [EG] is a source of potential threat to the [UJ] principles, questioning the tenacity of the foundational links between meanings and extensions, investigating its proper range is of utmost import for this ongoing project.

[§]

I now turn to some of the most important refinements, further developing this idea of an inverse relation between deference and transparency, before turning back to the question of the limits of externalism.

First, it is evident that there are many different kinds of use of any particular term; and so it is important to point out that the above remarks are premised on certain conceptions about typical uses by typical speakers. For example, notoriously, many uses of the term ‘water’ are not uses as a natural kind term, as above characterized. To delve further into this issue, consider

Putnam's (1962) thought experiment which concerns the surprising discovery that cats are actually Martian robots. Putnam (1962: 661) points out that intuitions may be divided between two different reactions to this surprising discovery:

- i. Wow! It's turned out that cats are not animals after all!
- ii. Wow! It's turned out that there aren't and never were any cats!

One fundamental difference between [i] and [ii] concerns exactly how the term 'cat' is used. Statement [i] involves a deferential, whatever-it-is-exactly-that-constitutes-the-real-essence-of-this, use of the term; whereas [ii] involves a more autonomous, my-meaning-determines-my-extension, use of the term. (The externalists' gap is open in [i], but not in [ii].) In my idiolect, cat is a [i]-type, natural kind term; and so [i] would be the correct response, while [ii] is confused.

Of course, someone else might insist on using 'cat' more autonomously, and insist that [ii] is correct of their use of the term. (Cf. Loar [1991: 120]: 'Social meanings do not deprive me of autonomy when I insist on it.' Deference cannot be forced, or legislated, onto a stubborn speaker.) This will afford more transparency and autonomy to their 'cat' thoughts and utterances. Fair enough; but, given that such stubborn speakers are not using 'cat' as a natural kind term, they pose no counterexample to idea that there is an inverse proportion of transparency and deference. The gains in transparency here come proportional to the lack of deference. However, such stubborn speakers do complicate any simple articulation of the inverse proportion relation, which does not take into account that any term can be used in multiple sorts of ways, according to the speaker's intentions.¹⁸

Again it is important to ward off oversimplistic misimpressions, caused by focusing exclusively on paradigm cases at the extreme poles. These paradigm cases are dialectically important; but clearly the lexicon embodies many complex and heterogeneous divisions, which map onto each other in complex ways, between deferential uses of 'aluminum' and transparent uses of 'and.' There may even be lots of cases which, upon reflection, we would want to classify in purgatory, between transparency and deference. For example, consider Putnam's (1962) 'Pencils are artifacts.' How would I respond to the discovery that pencils are actually intelligent organic Martian spies? Is this a discovery that there aren't and never were any pencils, or a big surprise about pencils? I for one have no strong intuitions either way on this one; which is to say that I use 'pencil' as neither a paradigmatically deferential natural kind term nor as a paradigmatically autonomous and transparent sort of term. (Cf. Putnam's [1975: 248] remarks about differences between 'cat' and 'pencil'; they come up again in the next section.)

Both of these refinements (i.e., that pertaining to different types of usage, and that pertaining to different types of term) show that we are clearly dealing with a complex variety of differences of degree along a continuum, as opposed to any straightforward, categorical difference in kind, from deferential uses of ‘aluminum’ to autonomous uses of ‘widow’ or ‘and’ (or: between natural kind terms on to social kind terms and logical particles). Further, there are experts who are not generally deferential about a given range of natural kind terms; and there are plenty of non-natural-kind technical terms about which most speakers are rather deferential (e.g., ‘hedge fund,’ ‘spandrel,’ ‘junta’). The picture is messy, as any picture of the semantic aspects of our linguistic behavior must be, if it aims at comprehensiveness. Nonetheless, if the broad contours of the picture traced herein are accurate, this has considerable significance for questions about the range of semantic externalism.¹⁹

To summarize and consolidate some of the points made about transparency and deference, then, consider the following five statements:

1. Aluminum is a metal.
2. Cats are animals.
3. Pencils are artifacts.
4. Widows are formerly married women whose spouse has died.
5. If P and Q, then P.

They are ranked in increasing order of resistance to correction. (The difference between [1] and [2], I take it, is that the aforementioned stubborn, non-deferential, non-natural kind uses of ‘cat,’ like of ‘water’ [cf. note 18], are more prevalent than such uses of ‘aluminum.’ No one would say that there isn’t and never was any aluminum!) They illustrate the inverse proportion of semantic transparency and semantic deference, in that each successive step brings with it less deference and more transparency. Furthermore, with each successive step, there is less and less purchase for the [EG] between what is constitutive of competence and what determines the extension.

§6.5: THE LIMITS OF EXTERNALISM

There are multiple different ways in which semantic externalism might be limited. For starters, we might distinguish two among them: limiting the RANGE of externalism versus limiting the DEPTH of externalism. For example, to hold that the seminal externalist arguments apply only to proper names and natural kind terms, and have no obvious relevance beyond those types of referring expression, is to limit the RANGE of externalism. In contrast, to hold that, even in the case of incompletely mastered referring

expressions, there are still significant limits to the amount of ignorance or mistakes which are compatible with competence, is to limit the DEPTH of externalism. (Surely even the most radical externalist has to agree that someone who thinks that ketchup is a kind of marsupial cannot be counted as competent with the term!) However, not much work has been done, when it comes to serious investigation of the DEPTH of semantic externalism—that is, defining minimum thresholds for incomplete-mastery-yet-competent.²⁰ Perhaps the only possible answer to that is stipulative. In any case, that is one pertinent sense in which semantic externalism is limited—no one is an externalist all the way down, holding that there are no substantive requirements for competence whatsoever, even in the case of deferentially used (Millian-tag-like) referring expressions.

I am presently concerned with the RANGE question: that is, How far, beyond the bounds of deferential uses of proper names and natural kind terms, can or should the seminal externalist arguments be extended? I will not touch on the DEPTH questions except in passing—not to mention on some other significant questions in the neighborhood, such as how externalism relates to the rule-following considerations, or to self-knowledge, or whether externalism entails *a priori* access to things one otherwise would have thought to be *a posteriori* (or vice versa), etc.

Recall from §5.3 the following three theses about the [EG]: the Pragmatic thesis takes the [EG] to be a property of certain limited type of linguistic *usage*; the Semantic thesis takes the [EG] to be a property of certain limited type of linguistic *expression*; and the Metasemantic thesis takes the [EG] to be a general discovery about the nature of *language*. I turn next to adjudicating this issue, and connecting it to our ongoing discussion.

For starters, then, why might one conclude, with the Pragmatic thesis, that the seminal externalist arguments tap into something distinctive about certain sorts of uses of a limited class of terms, as opposed to something about semantic competence in general? Well, the dialectic in Kripke's seminal work goes: (i) recent developments in modal semantics strongly suggest that proper names are much more like Mill thought them to be than like Frege or Russell thought them to be, (ii) and lo! They also show that natural kind terms are in many respects (i.e., those described above in §6.4) similar to names. That is, Kripke is quite plausibly read as arguing that proper names and natural kind terms (in typical sorts of usage) are *distinctively* non-descriptive. (Bach's [1987] 'relational determination' and Recanati's [1993] 'psychological neutrality' are some influential ways to further articulate what is distinctive about the mechanisms of reference involved in this circumscribed class of non-descriptive terms, ways that do not [at least not immediately, or obviously] extend across to linguistic meaning generally.)²¹

So, there are some positive reasons to think that the seminal externalists' arguments have tapped into something unique about a specific, circumscribed kind of referring expression, and hence to think that minimal conditions for competence apply distinctively to proper names and natural kind terms. (Again, this is most starkly evident in their deferential uses by non-expert speakers.) Insofar as 'Einstein' or 'tiger' is merely a tag or label—in a way that some if not most sorts of linguistic expression are not—then (i) the minimal threshold for competence, in such cases, is appropriately thought of as relatively insubstantial and undemanding, and (ii) competent grasp of such expressions will not bring in its train transparent access to much of anything. But this consideration does not yet apply to expressions which are not so plausibly thought of as mere tags or labels.

Note that Kripke himself enters some remarks contrary to Metasemantic thesis' range-extension, in the course of discussing the famous case of Pierre and 'London'/'Londres':

Not that the puzzle extends to all translations from English to French. At the moment, at least, it seems to me that Pierre, if he learns French and English separately, without learning any translation manual between them, *must* conclude, if he reflects enough, that 'doctor' and 'medecin,' and 'heureux' and 'happy' are synonymous, or at any rate, coextensive; and potential paradox of the present kind for these word pairs is blocked. (1979: 256)

(Not to suggest that Kripke is here explicitly concerned, let alone solely concerned, with the question of the range of the EG. However, there is surely a relevant distinction articulated here—albeit tentatively—about differential criteria for competence, which relates to transparency and externalism.) The thought here is that apart from the externalist-friendly (distinctively non-descriptive) cases of proper names (e.g., 'London') and natural kind terms (e.g., 'furze'/'gorse'), more substantive requirements are appropriate for competence.

Also, in the middle of the third *Naming and Necessity* lecture—Kripke's most extensive treatment of the [EG] for the case of natural kind terms—he expresses a similar tentative qualification to its range, citing 'foolish,' 'fat,' and 'yellow' (1972: 127) as examples of sorts of expression to which the ongoing discussion may not apply. So Kripke himself holds that it would not be easy, or obviously appropriate, to motivate an externalists' gap for the likes of 'foolish,' 'fat,' 'yellow,' 'doctor' or 'happy.' Hence, it is plausible to read Kripke as holding that many common expressions are descriptive. Further, similar remarks apply to Putnam too. He holds (1975: 233) that 'some words do not exhibit any division of linguistic labour: "chair," for example.' Further, he also concedes (1975: 248) that the case for an [EG] with respect to 'cat' 'has more plausibility' than the analogous case with respect to 'pencil'.

So, there do exist grounds for reading some of the great seminal externalists as moderate externalists, when it comes to the crucial question of differential criteria for competence, and for the relation between what constitutes competence and what determines reference. These big-picture considerations and historical precedents are confluent with the above cases considered, which run contrary to the Metasemantic thesis that the [EG] applies to all corners of the seas of language.

[§]

The externalist challenge, recall, is the consideration that factors external to individual speakers play various roles in determining the extensions of the terms which they employ in their thought and talk. It is most deeply relevant to charting *a priori*, because of its potential to undermine, or at least drastically alter, traditional conceptions of the relations between meaning and extension. The knock-on effects on the [UJ] connections would be immediate and drastic. Hence any radical form of semantic externalism would obviously affect our considered conception of (at least) semantic and epistemic immunity to counterexample (cf., e.g., Williamson [2008], Russell [2012] for concrete integuments of this line of thought).

Reasons are emerging, though, to hold that the Metasemantic thesis about the [EG] looks to be an unsupported RANGE-extension, when it comes to understanding the upshot of the seminal externalist arguments. There are some good reasons to hold that those arguments tap into something distinctive about a limited class of expressions; and there are some fairly strong counterexamples to radical externalist claims—as they pertain to, say, ‘grandmother,’ ‘square,’ and ‘and.’ In such cases, meaning limns extension, and so U can afford J.

What the counterexamples to semantic internalism (i.e., semantic content supervenes on the intrinsic states of individual speakers) clearly show is that otherwise competent speakers are often deferential as to the criteria for the correct application of the terms tokened in their thought and talk. In such cases, content is generally not transparent to individual speakers, and the externalist gap is evident between what is constitutive of competence and what determines the extension. However, there are vast stretches of the lexicon regarding which typical speakers are not typically (nor ought they to be) deferential about the criteria for correct application. These cases are beyond the scope of the seminal externalist arguments, per se, and in these cases there is no [EG] and content is transparent to competent speakers. As Burge (1979) points out, and we have further explored, degree of resistance to correction is a primary determinant for where a particular use of a particular term belongs on this deference-transparency spectrum.

In general, there is an inverse proportion of semantic transparency and semantic deference—that is, transparency wanes to the extent or proportion

that deference waxes. To the extent that the speaker is deferential in their use of a term, and is open to correction, then they lack transparent access to the conditions for the correct applicability of their terms. In contrast, to the extent that a speaker is not deferential in their use of a term, and is not open to correction, then the speaker does have transparent access to the conditions which determine the term's extension. (These non-deferential speakers are liable to make mistakes—incompetent moves—in a way that their deferential counterparts are not.) Clearly, a wide and complex range of cases stretches in between the extremes (illustrated herein by deferentially used natural kind terms and transparently used logical constants). Still, this might be seen to go some way toward undermining the Metasemantic thesis, which would be hostile to the [UJ] connections. The seminal externalist arguments have not shown the transparency of semantic content to be obsolete, beyond the bounds of deferentially used proper names and natural kind terms. It has not been shown that, in general and unequivocally, there is a lot less to semantic competence than the traditional internalists presumed. Incomplete mastery can amount to semantic incompetence, in lots of mundane cases, wherein deference is absent.

I next consider two questions which have come up:

[1] How far, beyond the range of referring expressions such as proper names and natural kind terms, does this externalist challenge extend?

The answer here depends upon the speaker's intentions—any term (even 'and') could be used deferentially. However, though, lots of terms are not in fact typically so-used. Furthermore, in any case, there are reasons to think that incomplete mastery is uniquely compatible with competence for the circumscribed cases of deferentially used proper names and natural kind terms, in a way that does not extrapolate across to 'pencil,' 'widow,' 'and,' etc.—because of certain distinctive properties of that circumscribed category of referring expressions. Insofar as 'Einstein' or 'tiger' is merely a tag or label—in a way that 'pencil,' 'widow,' 'and,' etc., are not—then the minimal threshold for competence, in such circumscribed cases, is appropriately thought of as relatively insubstantial and undemanding.

[2] Has it been proven that there is, quite generally, a lot less to semantic competence than traditional internalists would have us believe?

No, given that competent speakers are not always, and perhaps not even typically, deferential. (After all, I'll bet that in your house, like mine, there are a lot more chairs than tigers) Incomplete mastery is compatible with competence just to the extent that the speaker is deferential; this will differ

significantly (in degree) for distinct types of term, as well as for distinct types of usage for any given term. In any case, neither the seminal externalist arguments nor subsequent externalist work suffice to ground the need for any drastic, general re-conceptions of semantic competence. In lots of mundane cases, transparency holds, and (if the speaker in question is not appropriately deferential) incomplete mastery can still amount to semantic incompetence.

To sum this up, then: To endorse the Pragmatic thesis about the [EG] is one way to be a moderate externalist, conceding a main thrust of the externalist arguments but being discriminating about the range of the terrain they impact. All things considered, it is my preferred response to the externalist challenge—that is, there is an [EG], but it is a function of how certain limited kinds of terms are typically used by ordinary speakers. (Semantic externalism may nail shut the coffin lid on the very idea of *a priori* access to the essence of aluminum, but [E], widows, and fortnights remain steadfastly unaffected.) Furthermore, a moderate externalist response to the externalist challenge dovetails with a modal revisionist response to challenge of revisability, constituting a stable ground between absolutism and skepticism. Hence, the [UJ] connections can thereby survive these two formidable and important challenges, within a constitutive *a priori* orientation.

§6.6: SUMMARY OF PART III

Next, I summarize what the constitutive *a priori* orientation is, in general, and how it absorbs the challenges of revisability and externalism, in particular. Thus continues the ongoing process of charting my considered stance on the different senses of immunity to counterexample—mindful as always of the framework-relativity of immunity to counterexample, the context-relativity of frameworks themselves, and the important differences between different kinds of frameworks.

The most general question, over here in the corner of epistemology which overlaps with the philosophy of language and the metaphysics of modality is: Is there or is there not a non-empirical source of justification? Can an adequate epistemology be developed without appeal to one? *A priorists* are those who answer YES THERE IS to the first and NO IT CAN'T to the second. Some variants of classical skepticism and contemporary naturalism are two varieties of anti-*a priorist*; but probably most and in any case very many important philosophers are in the *a priorist* camp.

Historically, over the almost twenty-five centuries between Plato's 'light of the mind in her own clearness' and Bonjour's (1998) 'defense of pure reason,' most *a priorists* have espoused a variant of what was first distinguished in §1.4 above as an acquaintance-based approach, according to which

a priori is taken to be a quasi-perceptual relation between minds and the objects of knowledge, often described as a kind of ‘mental seeing’. The alternative understanding-based approach was perhaps first clearly articulated by Hume (1748) under the guise of ‘relations of ideas’, though one can clearly discern seeds of this approach in Hobbes and Locke, and perhaps going all the way back to Aristotle’s opposition to Plato’s extravagant, otherworldly epistemology. The guiding idea here is that the *a priori* is not an extra, elaborate cognitive faculty, but can rather be fashioned out of ingredients to which we are all committed (i.e., a comprehensive account of what it takes to grasp or understand a concept). The signature objection to acquaintance-based approaches is that they amount to little more than obscure, unhelpful metaphors; and a key challenge to understanding-based accounts is to explain the possibility of ampliative, substantive *a priori* knowledge. Critics of the latter allege that what it calls justification does not really amount to any such thing.

There are many significantly distinct sub-varieties of both acquaintance-based and understanding-based approaches to *a priori*. For example, even after narrowing our focus onto the understanding-based orientation, there are adherents to this view who take themselves to be rationalists, and others whose approach is styled in the empiricist tradition.

Following the lead of some major twentieth-century thinkers and developments, I am defending a constitutive *a priori* approach to this nexus at which epistemology overlaps with the philosophy of language and the metaphysics of modality. This is a broadly neo-Kantian status-refinement within the understanding-based tradition—that is, not only does *a priori* not apply to some special faculty of mind, as on the understanding-based tradition, it also does not solely apply to some special category of content, as the moderate rationalist would have it. It essentially also has to do with status, role, or function. To categorize something as *a priori* is, in part, to say something about its place in a given framework (or language game, world picture, theory, etc.).

I take the challenge of revisability, associated with Quine circa 1950, and the externalist challenge, associated with Kripke circa 1970, to be huge recent jolts to the *a priorists*’ world-order. They render if not obsolete, then at the very least in need of extensive refinement, many traditional ideas about epistemology, semantics, and metaphysical modality. I have argued that these challenges are well met by a constitutive *a priori* picture which espouses a modal revisionist response to revisability (i.e., the *a priori* is real but framework-relative) and a moderate externalist response to the externalists’ challenge (i.e., the [EG] is very real and deeply significant but limited in scope to certain types of uses of certain types of terms). This leaves a fairly broad and significant terrain open for business for acquaintance-based *a priori* ratiocination, amenable to [UJ]-underwritten immunity to counterexample.

Recall from §5.1 the following [UJ] principles, which I characterize as the spine of an understanding-based approach to *a priori*, and which articulate the direct contact which all of these finer points have to Plato's problem:

$$UJP1:[UJ \rightarrow (A \rightarrow AP)]$$

$$UJP2:[UBS \leftrightarrow (SI \rightarrow UJ)]$$

UJP1 tells us that if there are [UJ] connections, then analyticity can ground (at least some cases of) *a priori* justification. UJP2 tells us that the realm of the un-black-swannable is the range of cases in which semantic intuition can underwrite the [UJ] connections. Basically, to the extent that the [UJ] connections hold up, semantic intuition can be relied on to accomplish what has been traditionally demanded of rational intuition. In a vast range of cases, semantic intuition can underwrite immunity to counterexample.

Natural reality is the range of the necessary *a posteriori*; social-conventional reality is the range over which sense determines reference. It is here that the [UJ] connections are resistant to the [EG]. Again, though, bear in mind that these categories are not discrete monoliths. There are many shades of grey, and directions of fit, and multiplicities of overlapping frameworks, between deferential uses of 'Aluminum is a metal' and transparent and autonomous uses of 'A fortnight is a period of 14 days,' or 'If P and Q then P.'

Ultimately that is why my constitutive *a priori* view is a variety of empiricism, not rationalism—that is, the range of the *a priori* is the framework, and the concepts they compose; so there is no *a priori* knowledge of the world on this conception of this orientation. This I see as a natural, though perhaps not inevitable, consequence of the move to an understanding-based view, which denies that *a priori* is a distinct faculty of mind. It is starkly put by the externalist challenge, but the writing has been on the wall at least since the development of non-Euclidean geometries. Traditional conceptions of rational intuition are not the potent weapons they were once conceived to be.

In any case, the main thesis of Part III is that the constitutive *a priori* orientation on Plato's problem, and the framework-relative [UJ] connections which form its spine, can survive the challenge of revisability by endorsing modal revisionism, and the externalist challenge by endorsing a moderate variant of semantic externalism.

NOTES

1. Wittgenstein scholarship is a notoriously tangley, contentious, dangerous business—his work is subtle, nuanced, complex, and not to be read as supporting any

simple, straightforward philosophical thesis—so I enter the disclaimer that I have no pretense of having discovered the one true reading of Wittgenstein on the *a priori*. Rather, the following is a plausible reading of a certain thread running through Wittgenstein's thought, backed up with some textual evidence, chosen because I take it to well-serve my ongoing project. Not all experts will agree with all of it.

2. This may strike an echo to Poincaré's conventionalism: 'Although the influences on Wittgenstein's work are notoriously difficult to interpret, there are clearly evident similarities to Poincaré's fundamental views' (Heinzmann & Stump, 2013, §5). Relatedly, cf. Friedman (2007: 96–99) for conceptual ties between Poincaré and Reichenbach. There is a bit more discussion of Poincaré below.

3. For a sample of the literature on this issue cf. O'Grady (2004).

4. Wittgenstein wrote in a letter to Schlick in 1932 that 'neither Poincaré nor Reichenbach could have the same conception [of hypotheses as I do] because they do not share my conceptions of propositions and grammar'. Cf. Stern (2007: 325) for discussion.

5. As mentioned in the Preface, the constitutive *a priori* orientation has deep roots in the philosophy of science literature, from canonical European figures like Poincaré and Cassirer, to classical American pragmatists Pierce, Dewey, and Lewis. My aim here is not so much to exhaustively cover that historical terrain as to argue for its more broad applicability to the ways in which twentieth-century philosophy of language has altered traditional epistemology.

6. Friedman (2000) is extensively cited here as a representative illustration of the constitutive *a priori*, not because I think it is the final truth on the matter. Friedman's view continue to change and evolve (e.g., [2007, 2008, 2011], and so on), and there has been much critical discussion of Friedman's views which I will not turn to here. (For a variety of critical engagements with Friedman's work, cf. Korkut [2011], Angeloni [2012], Uebel [2012], Everett [2015], Stump [2015].)

7. Cf. Haack (1974) for a canonical discussion of deviant logics; cf. Beall & Restall (2006) for a case in favor of logical pluralism.

8. Williamson's version of what I am calling [UJLT corollary] is: if something is a logical truth, then assenting to it is a necessary condition for understanding it.

9. Actually, the ambiguity option might have some real purchase in this case. One could argue that ' $\sim\Phi$ ' or ' $\Phi \vee \Psi$ ' literally means something different for an intuitionist, as opposed to a classical logician. (Thanks to Wayne Myrvold for pressing this case in the discussion period after I presented the argument of this section.) To the extent that this is so, then this points to a rather clear difference between the DLOs of Horwich and of Williamson. Note well though that taking this path would undermine the promise of Horwich's DLO to support any drastic conclusions about [UJ] connections (which could only be derived from a SHARED understanding, not from diverse understandings).

10. As Flanagan (2013: 346–47) documents, multiple authors have tried to answer Williamson's challenge in one (or both) of these ways, but in Flanagan's (and my) assessment, that will not do. As Williamson (2009: 135; 2011: 499) insists, this is a case of 'theoretical disagreement,' not equivocation or incompetence.

11. Williamson (2008: 95) concedes that [&E] may 'have the best chance' as far as candidates for [UJ] connections go, but argues that even it is subject to competent, unequivocal dissent. However, his putative counterexamples (p. 96) are relatively weak

and problematic. (Cf. Boghossian [2011], Peacocke [2011] for discussion.) In my opinion he does a much better job of motivating counter-instances for ‘All As are As’, and they are enough to force a challenge to [UJLT].

12. Indeed, within propositional logic, one might take the LEM and LNC to implicitly define the term ‘proposition’. (The price of that move is that ‘There will be a sea battle tomorrow’ or ‘Erin is tall’ might fail to express a proposition.)

13. Maddy (2000: 113) reads Carnap in a way that is incompatible with this kind of complexity; I will not get into a battle of Carnapian exegesis but will explicitly disavow any such static, categorical reading of the internal/external or language/theory distinctions as applied to frameworks.

14. Railton’s messy Wittgensteinian picture stands in stark contrast with Friedman’s much more regimented three-level picture. This is due in no small part to differences of range—that is, Friedman is focused on the history of science, while Railton is evidently talking about inquiry in general.

15. And this is not even to mention her huge qualification about the case of logic in her note 43. If we compare the quote from Quine (1951: §6) cited at the opening of §5.1, about revisability within logic, there are some serious doctrinal differences between Maddy’s naturalism and Quine’s.

16. As alluded to above in §5.4, Locke’s (1690: Bk III) distinction between terms which target a real essence and terms which target a nominal essence is an important early analysis of this core distinction. That widows are formerly married women whose spouse has died articulates a nominal essence, while that aluminum is a metal targets a real essence. Deference, openness to correction by experts, and the possibility of community-wide error are distinctively evident in cases of intending to target real essence.

17. Here we brush up against the issue of tacit knowledge, a general issue in the philosophy of psychology concerning the different senses in which (say) competent speakers should be said to ‘know’ the rules which govern their linguistic behavior. For example, anyone who has taught Introductory Logic will have encountered students who unflinchingly use ‘and’ in accordance with the standard introduction and elimination rules, but who nonetheless have great difficulty in attaining a theoretical grasp of those rules. Such agents tacitly know these rules, as is attested by their (deliberate, rule-governed, intentional) ground-level thought and talk involving ‘and’, even despite their theoretical difficulties.

I take it that such agents are not counterexamples to the idea that competence with ‘and’ is constitutively tied to the standard introduction and elimination rules—only someone who we would want to call competent with ‘and’ but yet reflectively conceded counterexamples to those rules would. (Cf. §6.2 for extensive discussion of such deviants.) However, they do show that finessing is required as to exactly what is meant by ‘transparency’ in this context. Such agents show that my claim has to be that grasp of intension can be transparent while being merely tacit, as opposed to the stronger claim (which I do not want to make) that transparency demands the kind of theoretical grasp which is necessary for success in a Logic course. What matters for the claims about ‘transparency’ made herein is (deliberate, rule-governed, intentional) action in accord with the rules, and that such rules are in principle accessible to reflective agents—despite the evident non-trivial differences between competent speakers, when it comes to the ability to consciously grasp the rules.

The general issue of tacit knowledge has some relevance to this ongoing strand of debate between the different theses about the [EG] introduced in §5.3 and returned to immediately below, insofar as they are essentially (if only partly) concerned with what counts as semantic competence. However, this relevance is relatively tangential—it is not as if the extreme differences between Locke (1690) and Soames (2002), when it comes to competence, have anything to do with tacitness per se—and so I will go no further into it here.

There is a bit more on Williamson's challenge to the common presumption that competence with 'and' is constitutively tied to the standard introduction and elimination rules below at note 19.

18. We need not even go to such far-fetched thought experiments to illustrate this point. For example, reactions [i] and [ii] might surely have applied to the scientific recognition that whales are mammals and not fish, or, more recently, to the de-classification of Pluto as a planet. (This is a nice recent illustrative example: media reports I encountered about the decision to de-classify Pluto from the ranks of the planets tended to include a stubborn person-on-the-street insisting 'As far as I am concerned, Pluto will always be a planet!' This is a transparent but not deferential response.)

This first refinement is also applicable to the Chomskyan (1993, among other places) sort of pessimism about the externalist tenet that H_2O is the essence of water, on the grounds that what counts as water in lots of places is a lot less purely H_2O than Sprite or tea is. (Cf. note 11 from chapter 5.) Many such uses of 'water' are not uses as deferential, essence-targeting natural kind terms, but rather uses as relatively crude practical kind terms (i.e., 'whatever it is that flows out of this tap').

19. This point is relevant to debates between Boghossian and Williamson, over the [UJ] sort of epistemology of logic I am appealing to for 'and' (i.e., the presumption that competence with 'and' is constitutively tied to the standard introduction and elimination rules). One of Williamson's challenges to Boghossian's defense of this epistemology of logic is to say: if all your epistemology can guarantee is such pristine and sterile instances like 'and'-elimination, then your epistemology is not terribly worthwhile:

Strategically, Boghossian's response is not very promising. If he can rely on understanding-assent links only for 'and'-elimination and a few other equally banal rules ... then he is in no position to base either a *general* epistemology of logic or a *general* account of the understanding of logical constants on understanding-assent links. (2011: 498)

I see in this inverse proportion relation between transparency and deference the promise for a general reply to Williamson on this front, which applies to varying degrees, as appropriate, beyond the bounds of 'and'. (Of course, I am not sure that Boghossian would want to endorse anything along these lines.)

20. There is of course lots of relevant literature, on similar and related questions—such as, for example, the rule-following considerations (e.g., What exactly constitutes meaning *addition* by '+')? My claim is just that direct attempts to explicitly investigate the depth of competence required in these externalists' incomplete-mastery cases are relatively rare. (Cf., Sullivan [2010] for discussion and references.)

21. This line of thought is extensively developed in Sullivan (2012).

Part IV

**MAPPING OUT A CONSTITUTIVE
A PRIORI VIEW**

Chapter 7

Entailments and Conclusion

§7.1: TWO MAPS OF THE TERRAIN

The aims of this final chapter are, first, to summarize the constitutive *a priori* orientation as an answer to Plato's problem, and, second, to consolidate and extend some main points about the nexus at which epistemology overlaps with both the philosophy of language and the metaphysics of modality.

At a bit more length: this chapter is structured around two related maps. The first map will be developed in §§7.2–4. It concerns the following table, which charts eight possible permutations of the three key notions of immunity to counterexample—necessity, analyticity, and *a priori*.

That is, for example, row #3 represents the circumstance that something might be metaphysically necessary and knowable *a priori*, and yet fail to be analytically true; row #5 represents the circumstance that something could be analytic and *a priori*, and yet fail to be necessary; and so on.

Table 7.1

1	N+	A+	AP+
2	N+	A+	AP–
3	N+	A–	AP+
4	N+	A–	AP–
5	N–	A+	AP+
6	N–	A+	AP–
7	N–	A–	AP+
8	N–	A–	AP–

Rows #1 and #8 should be the least controversial, and the most familiar, roughly corresponding to Hume's 'Relations of Ideas' and 'Matters of Fact,' respectively. (Or, if you like, Plato's Parmenidean side and Plato's Heraclitean side, respectively.) The strongest candidates for row 1 would include fundamental truths of mathematics and logic, such as [1] and [2]:

1. No proposition is at once both true and not true.
2. Two is a factor of every even number.

Again, even these are not completely unanimous, as paraconsistent logics reject [1], and certain varieties of constructivism may balk at the claim that there is anything metaphysically necessary about [2]. Still, the claim that these belong in row 1 is rather orthodox, probably as close as one can get to unanimity in philosophy. For present purposes, I am content with the qualified claim that these are the strongest candidates, and orthodox picks, for row 1, and avoid the arguments and counterarguments on this point.¹ As for row 8, candidates are mundane and ubiquitous—for example, 'There are ten provinces in Canada,' 'Jupiter has more moons than Neptune,' etc.

What, then, of the other six rows? The controversies over those regions, as they look through the lens of a constitutive *a priori* orientation, will occupy us for the next three sections.

These issues will be approached down a slightly different avenue in §7.5. Here the discussion will be structured around evaluating the following six conditionals:

1. $N \rightarrow A$
2. $N \rightarrow AP$
3. $A \rightarrow N$
4. $A \rightarrow AP$
5. $AP \rightarrow N$
6. $AP \rightarrow A$

That is, to endorse [1] is to hold that that all metaphysical necessities are analytic truths; to endorse [6] is to assert that everything knowable *a priori* is analytically true; and so on. The aim of this part of our investigation is to consolidate what has been learned about the plausibility of any such conditionals, especially in the wake of Quine's challenge of revisability and Kripke's externalist challenge.

Finally, §7.6 will consist of summarizing our results, drawing out some generalizations, and looking back over some of the perennial themes which have woven through the book.

§7.2: CAN ANALYTICITY VARY INDEPENDENTLY?

Getting back to our first map, then:

Table 7.2

1	N+	A+	AP+
2	N+	A+	AP-
3	N+	A-	AP+
4	N+	A-	AP-
5	N-	A+	AP+
6	N-	A+	AP-
7	N-	A-	AP+
8	N-	A-	AP-

Perhaps the least likely contenders on the table are the rows which allow analyticity to vary independently—that is, 3 [N+, A-, AP+] and 6 [N-, A+, AP-]. Could there be an analytic truth that is neither necessary nor knowable *a priori*? In the other direction, could something be necessarily true and knowable *a priori*, and yet fail to be an analytic truth?

Let us take row 6 first. The very idea of an analytic truth that is neither knowable *a priori* nor necessary seems jarring, to say the least—perhaps, again, due to the distinctive geography of the semantic, constituting as it does a bridge between mind and world. Analyticity, it seems, cannot stand on its own, but rather needs to be co-instanced with (and further, perhaps, grounded in) at least one of the requisite metaphysical furniture or the requisite epistemic relations. To deny that row 6 has any denizens is to assert the following conditional:

[A → AP ∨ N] If something is an analytic truth, then it is either knowable *a priori* or necessarily true.

I can think of no counterexamples to this, nor of any philosopher who has even implicitly transgressed this conditional.

In contrast, row 3 [N+, A-, AP+] does have some historical precedent. In particular, (given Kant's presumption that necessity is a criterion for *a priori*) Kant's putative synthetic *a priori* judgments might be thought to belong in row 3. If so, then the fate of row 3 may be tied to some rather huge issues, within (and well beyond) epistemology. (For example, as cited in §1.4 above, both BonJour [1998] and Aune [2008] hold that the fate of rationalism itself rests on the coherence of the synthetic *a priori*.)

However, among the major problems to be ironed out before we should concede that certain Kantian candidates belong in row 3 is that Kant gives no

firm and comprehensive criterion for analyticity.² Further, I believe that many of Kant's specific candidates for synthetic *a priori* status have been more or less conclusively refuted. (I have in mind here especially the mathematical and logical examples, falling prey to the careful distinctions and counterarguments of especially Frege.) To be sure, the same cannot be said of all of Kant's candidates, such as the following:

1. Every event has a cause.

(Assume for present purposes that [1] is true, and is known to be true.³) In my view, this is among the most plausible candidates for synthetic *a priori* status. It seems at once not obviously constitutive of the concept 'event'—that is, an uncaused event would hardly be a contradiction, as would, say, a five-sided square—and hence it is synthetic; while at the same time there is some inclination to hold that one's grounds for believing [1] are non-empirical. (It feels more solid than a typical black-swannable inductive generalization.)

However, I take the decisive problem at this juncture to be Kant's presumption that necessity is a criterion for *a priori*. That is, the inclination to count [1] as *a priori* stems, in large measure, from the intuition of its necessity. If [1] is true, then there are grounds to deny that it could be so by contingent happenstance. It is plausible to think that, if it is true, what makes it true must be some deep, fixed, basic features of objective mind-independent reality. So, Kripke's (1971, 1972) guiding intuition of the form '*IF P is true, then it is necessarily true,*' which plays a key role in motivating the notion of the necessary *a posteriori*, seems applicable to [1].⁴ Given that, and since reasons have subsequently emerged to doubt that necessity is a criterion for a *priority*, Kant's case for the *a priori* of [1] is considerably weakened.

My own view is that [1] is, if true, a Kripkean necessary *a posteriori*. Justification is not merely non-empirical, in such a case, and so it belongs in row 4 [N+, A-, AP-], not in row 3 [N+, A-, AP+]. (These claims will be bolstered in the next section.)

The reasons to doubt Kant's presumption that necessity is a criterion for *a priori* leave us without a positive reason to populate row 3, now conjoined with a standing inclination to hold that analyticity cannot fall on its own, any more than it can stand on its own. If, accordingly, we close off row 3, then we assert the following conditional:

[AP & N→A] If something is both knowable *a priori* and necessary, then it is an analytic truth.

If both the requisite metaphysical furniture and the requisite epistemic relations are in place, then the relevant semantic relation is bound to be instanced. (See note 7 below for potential qualification.)

To take stock: Rows 1 and 8 are the least problematic, while rows 3 and 6 are quite problematic. That leaves rows 2, 4, 5, and 7 as open questions.

Table 7.3

1 ☺	N+	A+	AP+	Two is a factor of every even number.
2 ?	N+	A+	AP-	
3	N+	A-	AP+	AP & $N \rightarrow A$
4 ?	N+	A-	AP-	
5 ?	N-	A+	AP+	
6	N-	A+	AP-	$A \rightarrow AP \vee N$
7 ?	N-	A-	AP+	
8 ☺	N-	A-	AP-	Canada has ten provinces.

I will divide up these four remaining open rows into the Kripke cases (i.e., rows 2 & 4) and the indexical cases (i.e., rows 5 & 7).

§7.3: THE KRIPKE CASES

By ‘the Kripke cases’ I mean certain varieties of what Kripke (1971, 1972) argued should be understood as necessarily true but nonetheless only knowable *a posteriori*. Key here are the cases in which it is plausible to hold that the essence of a natural phenomenon has been scientifically discovered. Candidates include:

2. Heat is the motion of molecules.
3. Gold is the element with atomic number 79.

On the one hand, it is arguable that these are necessary truths, in that any possible phenomenon which satisfies the subject also and thereby satisfies the predicate, and vice versa. On the other hand, these can only be known *a posteriori*—armchair reflection on the subject-concepts will not suffice; but, rather, lots of empirical evidence is required.

Now, to be sure, there have been considerable objections to Kripke’s arguments, which I will not try to comprehensively address here.⁵ (To cite just one example, the argument in favor of their necessity relies on a premise of the form ‘*IF P is true, then it is necessarily true*’; that premise is surely *a priori*, which may go some way to undermine the intuition that this should be counted as *a posteriori* knowledge.) Further, many of Kripke’s putative candidates are more controversial

than [2]–[3].⁶ Nonetheless, to the extent that: (i) one of the goals of scientific inquiry is to discover the essences of natural phenomena, and (ii) scientific findings are, by and large, *a posteriori*, then there can be overlap between the categories of necessary truth and *a posteriori* knowledge. I take [2] and [3] to be fairly uncontentious cases, which are enough to show that at least one of either rows 2 or 4 is occupied.

So, then, do these cases belong in row 2 or 4? That is, are these necessary *a posteriori* truths analytic or synthetic? As ‘analyticity’ is defined above, [2] and [3] are clearly not analytic, because to judge them false may be mistaken, but is surely not a contradiction. (Heat without molecular agitation would be a significant and surprising discovery, but would be nothing along the lines of the discovery of a five-sided square!) Alternatively, one could be competent with their subject-terms without having any opinion as to the truth of [2] or [3], and so they are not true in virtue of the meanings of their constituent terms. So, [2]–[3] belong in row 4 [N+, A-, AP-]. In other words, the relevant extensions may stand in necessary relations; but, given the degree of deference appropriate to the use of a natural kind term, there is no transparent analytic or *a priori* access to their so-standing. (Again, though, as discussed in Part III, such cases might become analytic necessities, as a function of conceptual evolution. But they would thereby also become knowable *a priori*, and hence row 1 cases. Migration from row 4 to row 1 is possible, and may well happen over a few generations for a case like ‘Whales are mammals’.)

Row 2 [N+, A+, AP-], it seems, is bound to end up empty. It is, to say the least, hard to see how something could be both necessary (i.e., could not fail to be) and analytic (i.e., roughly, true in virtue of meaning) but not knowable *a priori* (i.e., roughly, justified non-empirically); and I cannot think of much in the way of plausible counterexamples nor historical precedent to the contrary.⁷ So, then, here is where the investigation of the Kripke cases leaves us:

Table 7.4

1 ☉	N+	A+	AP+	Two is a factor of every even number.
2	N+	A+	AP-	A & N→AP
3	N+	A-	AP+	AP & N→A
4 ☉	N+	A-	AP-	Heat is the motion of molecules.
5 ?	N-	A+	AP+	
6	N-	A+	AP-	A→AP ∨ N
7 ?	N-	A-	AP+	
8 ☉	N-	A-	AP-	Canada has ten provinces.

§7.4: THE INDEXICAL CASES

Now to what I am calling ‘the indexical cases’—that is, certain putative candidates for rows 5 and 7. While denizens of this terrain have played important roles in philosophy since at least Descartes (1641), it is arguable that there were no sophisticated maps of this area until the 1960s and 70s. During that period, there occurred the development of multi-dimensional logics in which the effects of the context of utterance on content semantically expressed can be neatly distinguished from the effects of the context of evaluation on whether that content is true or false.⁸ For example, consider the exact sense in which the following is immune to counterexample:

6. I am here now.

Following Kaplan (1989), among others, it is not acceptable to count this as a necessary truth. For any utterance of [6], the indexicals are saturated by relevant aspects of the context of utterance, and the content expressed is constitutively tied to a specific individual, place, and time (e.g., Arthur is in his kitchen at 7:00 am). Clearly, that expressed content (or any similarly specific expressed content) is contingent, for had accidents gone otherwise I could have been in my car or on a plane at that instant. At the same time, though, [6] is a good candidate for both analyticity and *a priori*. It is a good candidate for analyticity because it is precisely the semantic contents of its constituent bits that ensure that what it expresses is, even though contingent, nonetheless guaranteed to be actually true. It is a good candidate for *a priori* because the kind of justification one would give for one’s confidence that an utterance of [6] expresses a truth would not be empirical. (I would hardly have to investigate as to exactly where I was, before I felt justified in believing that [6] expresses a truth.) So, row 5 it is. [6], among some other cases,⁹ is [N-, A+, AP+]:

Table 7.5

1 ☉	N+	A+	AP+	Two is a factor of every even number.
2	N+	A+	AP-	A & N→AP
3	N+	A-	AP+	AP & N→A
4 ☉	N+	A-	AP-	Heat is the motion of molecules.
5 ☉	N-	A+	AP+	I am here now.
6	N-	A+	AP-	A→AP ∨ N
7 ?	N-	A-	AP+	
8 ☉	N-	A-	AP-	Canada has 10 provinces.

[§]

The last remaining open question, as to the fate of row 7, turns on how strictly we construe the term ‘experience’ in defining *a priori* knowledge. (Recall that these matters were delved into, in considerable depth, in §4.2 above.) Is the relevant sense of ‘experience’ limited to sensory perception, or ought non-sensory introspection to also count as experience? How one approaches that question will determine whether ‘I exist,’ or ‘I am conscious,’ or perhaps even ‘I am hungry,’ will turn out to be justifiable *a priori*. They are independent of perceptual experience, to be sure, but hardly independent of experience *tout court*, in every philosophically interesting sense of the term.

I will not get into the matter of trying to adjudicate which of these options has in its favor the greatest philosophical precedent. (To a large extent, the exciting action has not been focused on this question, due to such factors as: (i) insofar as perceptual beliefs are the paradigm for the *a posteriori* case, the question of whether there can be non-perceptual *a posteriori* beliefs does not arise, and (ii) insofar as, over a broad range of paradigm cases, the *a priori/a posteriori* distinction lines up neatly with the distinction between beliefs whose content is general [or universal, or necessary] and beliefs whose content is singular [or local, or contingent], the question of whether there can be *a priori* but singular [local, contingent] beliefs does not arise.) Instead, I will first describe the fate of row 7 [N-, A-, AP+] on a narrow construal of ‘experience,’ and then describe the situation on a broad construal of ‘experience.’

First, if we construe ‘experience’ narrowly, such that only perceptual experience counts, then a belief is justified *a priori* iff its justification is independent of perceptual experience. Second, if we go with a broader sense of ‘experience,’ then something is justified *a priori* iff its justification is independent of experience *tout court*. I will call the first, narrow sense ‘APⁿ,’ and the second, broad sense ‘AP^b.’ Thus, the likes of ‘I exist,’ ‘I am conscious,’ and perhaps even ‘I am hungry,’ count as APⁿ but not AP^b.

For the case of APⁿ, then, there are some fairly plausible candidates for row 7 [N-, A-, AP+], such as:

7. I am conscious.

[7] is clearly not a necessary truth. Yet it is APⁿ, for its justification is not constitutively tied to sensory evidence—for I could be justified in believing [7] while not receiving, or attending to, any perceptual inputs. However, unlike [6], it is not analytically true. The analyticity of [6] rests on the point that the meaning of ‘I’ plus the meaning of ‘am’ plus the meaning of ‘here’ plus the meaning of ‘now’ suffice to ensure that

a truth is expressed. No such point holds for the case of [7]. (I am, if you like, semantically guaranteed to be here now, but not semantically guaranteed to be conscious.) So, for the case of AP^n , the table looks as follows.

Table 7.6

1 ☉	N+	A+	AP+	Two is a factor of every even number.
2	N+	A+	AP-	A & $N \rightarrow AP$
3	N+	A-	AP+	AP & $N \rightarrow A$
4 ☉	N+	A-	AP-	Heat is the motion of molecules.
5 ☉	N-	A+	AP+	I am here now.
6	N-	A+	AP-	$A \rightarrow AP \vee N$
7 [AP^n] ☉	N-	A-	AP+	I am conscious.
8 ☉	N-	A-	AP-	Canada has ten provinces.

Here we have some Cartesian comfort for Kantians—that is, a sense in which there is synthetic *a priori* knowledge. (Cold comfort it is, though, because the link to necessity has been severed. Kantians would presumably say that it is just not the same.)

[§]

Turning now to the case of AP^b , it seems that the above sorts of case cannot arise. In this case, [7], along with ‘I exist,’ ‘I am hungry,’ etc. will come out as *a posteriori*, because the justification for such cases does constitutively depend on some relevant senses of ‘experience’ (even if it is independent of current perceptual experience). That is, the explanation for why [7] is immune to counterexample, say, ineliminably involves an appeal to subjective qualitative feel, to the characters or qualities of experiences (which need not be perceptual). This broad construal of experience yields a more symmetrical table.

Table 7.7

1 ☉	N+	A+	AP+	Two is a factor of every even number.
2	N+	A+	AP-	A & $N \rightarrow AP$
3	N+	A-	AP+	AP & $N \rightarrow A$
4 ☉	N+	A-	AP-	Heat is the motion of molecules.
5 ☉	N-	A+	AP+	I am here now.
6	N-	A+	AP-	$A \rightarrow AP \vee N$
7 [AP^b]	N-	A-	AP+	$AP^b \rightarrow A \vee N$
8 ☉	N-	A-	AP-	Canada has ten provinces.

AP^b , then, entails at least one of analyticity or necessity, while AP^n , in contrast, can stand on its own (e.g., ‘I am conscious’ is AP^n but not analytic). Thus, conjoining that point with earlier results, we get a biconditional between analyticity and AP^b :

$$A \leftrightarrow AP^b$$

but only a conditional between analyticity and AP^n :

$$A \rightarrow AP^n$$

Neither construal of *a priority*, though, is conditionally linked to necessity.

Whether you want to call that a consequence of Kant’s (1781) status-orientation on *a priority*, or of Kripke’s (1972) clear distinction between satisfying the conditions for *a priority* and for necessity (or both, perhaps among other things [e.g., Poincare, Reichenbach, Wittgenstein, Carnap, etc., might also be cited as having provided key pieces of this puzzle]), on the constitutivist picture there is no immediate inference back or forth between necessity and *a priority*.

§7.5: CONSOLIDATING THE ENTAILMENT RELATIONS

An alternative map of this same terrain, on a constitutive *a priori* view, will result from evaluating these six conditionals:

1. $N \rightarrow A$
2. $N \rightarrow AP$
3. $A \rightarrow N$
4. $A \rightarrow AP$
5. $AP \rightarrow N$
6. $AP \rightarrow A$

It is an important legacy of mid-twentieth-century philosophy—starting with the development of rigorous semantics for modal logics in the 1950s, and gaining strength with the detection of fallacies and confusions in several varieties of argument against the coherence of metaphysical necessity¹⁰—that [1] and [2] are not terribly plausible. Given even a rather minimal degree of metaphysical realism, it becomes awfully difficult to conclusively establish that something is a necessary truth, and no epistemic or semantic conclusions are directly entailed by the claim that something is a necessary truth. (Again, this point is driven home most starkly by the externalist challenge,

especially as it pertains to natural kind terms.) Metaphysical necessity lies beyond the control of thought or language. And so we are down to at most four conditionals:

1. $N \rightarrow A$
2. $N \rightarrow AP$
3. $A \rightarrow N$
4. $A \rightarrow AP$
5. $AP \rightarrow N$
6. $AP \rightarrow A$

It seems that [3] and [5] are also off the table. What I above refer to as ‘the indexical cases’ show that something can satisfy either the definition of ‘*a priori*’ (i.e., roughly, justifiable non-empirically) or of ‘analytic’ (i.e., roughly, true in virtue of meaning) without satisfying the definition of ‘necessary’ (i.e., roughly, could not possibly fail to be). This leaves us with at most two conditionals:

1. $N \rightarrow A$
2. $N \rightarrow AP$
3. $A \rightarrow N$
4. $A \rightarrow AP$
5. $AP \rightarrow N$
6. $AP \rightarrow A$

On an understanding-based orientation, [4] is the most firmly grounded of the six possible conditionals. Semantic intuition provides justification for at least a broad sub-class of what traditional rationalists have wanted out of rational intuition. All analytic truths are knowable *a priori*. Note well, though, that it *is* a conditional—that is, it has not really been a direct aim of the present work to conclusively counter all of the arguments against the intelligibility of the notion of analytic truth.¹¹ Nonetheless, if there is such a thing as analytic truth, the present considerations suggest that it will play an important role in any adequate map of the *a priori*.

As for the fate of [6], see the close of the last section—that depends on whether one takes the relevant notion of *a priority* to be AP^b or AP^n . AP^b entails analyticity, but AP^n can stand on its own.

§7.6: IMMUNITY TO COUNTEREXAMPLE AND PLATO’S PROBLEM

A point of departure for this inquiry is the intuition that one central thing that the concepts of necessity, analyticity, and *a priority* have in common—and

which, to a large extent, accounts for why they are of such deep and enduring philosophical interest—is that all are tightly linked, in some way or other, to what we might call *immunity to counterexample*. Throughout, we have distinguished between three different types of immunity to counterexample—one metaphysical, one semantic, and one epistemic. A proposition may be necessarily immune to counterexample in virtue of the nature of the mind-independent facts of the matter (e.g., no two solid objects can simultaneously occupy the same spatial location); a statement may be analytically immune to counterexample in virtue of the meanings of its constituent parts and the way in which they are arranged (e.g., no grandmothers are childless); a piece of knowledge may be *a priori* immune to counterexample in virtue of one's justification for believing it being independent of experience (e.g., one cannot steal one's own property).

The externalist challenge shows up a potential wedge between meaning and extension (at least for certain sorts of cases—especially natural kind terms in their typical deferential use by non-experts); relatedly, there is a gap between metaphysical necessity on the one hand, and the semantic and epistemic modalities on the other. The counterexamples to the traditional semantic internalist idea that meaning determines extension run parallel to the related point, on a different level, that no semantic or epistemic conclusions follow *per se* from the claim that a certain proposition is necessary.

It is also quite closely related that the challenge of revisability turns out to not apply to metaphysical modality, but to prompt important revisions to the semantic and epistemic cases.¹² The notion of the framework-relative constitutive *a priori* promises to be a good way to absorb the shocks to the modal world-order prompted by these two challenges, and to afford an adequate and non-obscure answer to Plato's perennial question about the contents of our knowledge outstripping the limits of our particular experiences.

As there exist deep constitutive links between the epistemic and semantic cases (between our beliefs and the meanings of which they are composed), similar points will also apply to the notion of analytic truth. Analyticity is also a framework-relative notion, and hence subject to revision over time in cases of conceptual evolution. Nonetheless, there remains on this stance a clear and substantive sense in which the status of *a priori*, and the property of analyticity, constitute a kind of immunity to counterexample. The meaning-to-extension links, and hence [UJ] connections, are still evident in a broad range of cases.

However, conceivability does not entail possibility, on this constitutive *a priori* orientation, because of the deep fissure separating the semantic and epistemic modalities, on the one hand, from metaphysical modality, on the other. Conceivability is a main source of evidence for possibility, but it can founder on coldly indifferent metaphysical rocks. Judgments of conceivability tell us about frameworks of meaning, not about extensions.

Necessity, analyticity, and *a priori* are one and all indispensable elements of the philosophers' toolkit. Some crucial differences between them which have emerged are key defining features of this constitutive *a priori* view.

[§]

Next then to a quick recap and summary. We have distinguished, in the course of our investigation, two distinct types of metaphysical necessity. There is the row 1 type (i.e., N+, A+, AP+), and the row 4 type (i.e., N+, A-, AP-). And note again that there seems to be some possibility for gradual migration from row 4 to row 1, in certain kinds of cases of conceptual evolution. For example, 'whales are mammals' was initially a row 4 case, and remained such for generations after its discovery. However, it may well become a row 1 case, provided that we reach a point at which we would want to say that anyone who doubts whether whales are mammals thereby does not share our concept of 'whale.' (Certainly, 'cats are animals' feels solidly row 1 to me—even if it is conceivable that cats turn out to be Martian spies, that doesn't make it possible! And so presumably 'whales are mammals' could get there. Both contrast with 'Water is H₂O,' which is a less plausible candidate for migration to row 4 because of its relative technicality.)

We have also distinguished two sub-varieties of analytic truth. In addition to the row 1 cases, there are also the row 5 cases (i.e., N-, A+, AP+). On this front, developments in the semantics of indexicality have shown that strong modal status on the semantic (and epistemic) front does not entail metaphysical immunity to counterexample. Significantly, on this understanding-based, constitutive *a priori* orientation, all analytic truths are knowable *a priori*. (Anything guaranteed by semantic intuition is already sufficiently well-grounded, before we need to get into difficult questions about rational intuition.)

Given the complexities attendant upon the different senses of 'experience' which might be pertinent to fleshing out 'independent from experience,' detailed above in §§5.2 and 7.4, *a priori* is the most complex case of the three for registering simple generalizations. There are, again, row 1 *a priorities* (i.e., N+, A+, AP+), as well as row 5 *a priorities* (i.e., N-, A+, AP+). Further, if we construe 'experience' narrowly, then there will also be row 7 *a priorities* (i.e., N-, A-, APⁿ+)—such as 'I exist' or 'I am conscious.'

As for the great Kant, the constitutive *a priori* orientation is in many deep ways indebted to his work. Kant offers one of the most original, insightful, and seminal answers to Plato's problem in the history of Western philosophy. However, many reasons for departing from the letter of Kant's doctrines have been detailed herein. My two principle objections to Kant's synthetic *a priori* are: (i) Kant's account of the analytic/synthetic distinction is deeply flawed, and (ii) that Kant's presumption that necessity is a criterion for *a priori*

is not warranted. As the discussion of the ‘Every event has a cause’ case in §7.2 illustrates, many of Kant’s putative synthetic *a priori* cases survive in this present map as Kripkean necessary *a posteriori* row 4 cases (i.e., N+, A-, AP-). Once we absorb the shocks to the modal world-order in the latter half of the twentieth century, that seems to be a more appropriate classification for the informative discoveries which Kant placed at the center of his philosophy.

[§]

I come next to some remarks in response to a claim made by both BonJour (1998) and Aune (2008), cited above—that is, that the fate of rationalism itself rests on the coherence of the synthetic *a priori*. I am initially inclined to disagree with this claim, on the grounds that the fate of rationalism rests more squarely on whether there is *a priori* knowledge. That would be enough to keep rationalists gainfully employed; they do not require the further claim that at least some of that *a priori* knowledge cannot be, in any sense, analytic. However, still, surely BonJour and Aune are onto something here—if, say, Hume (1748) or Ayer (1936) are right that all *a priori* knowledge is analytic, then in a fairly clear sense there is no work to be done by rational intuition. And what remains of rationalism, without rational intuition?

Well, one line of response is suggested by the above finding that APⁿ is distinct from analyticity. That is enough to prove that these concepts differ not only intensionally but extensionally, and hence that there can be substantive differences between approaching a claim via the question ‘Is this analytic?’ and approaching it via the claim ‘Is this *a priori*?’

More deeply, though, while the conjecture that all necessary truths are analytic feels more like an attempt to explain away, rather than to explain, metaphysical modality—at least arguably, embodying a reductive, positivistic meta-philosophy—in contrast, the conjecture that all analytic truths are knowable *a priori* is a substantive theory of the *a priori*. It need be no part of this latter view that *a priority* is an illusion, or that *a priority* has been systematically mis-categorized (as is, say, the idea that linguistic conventions are the real source of what we mistakenly take to be our intuitions of metaphysical necessity). This is, rather, an attempt to explain or to ground *a priority*.

So, whereas to appeal exclusively to meanings in answering metaphysical questions is a deeply problematic restriction, in contrast, meanings are integrally engaged with, and largely constitute, the data to be explained by epistemology. Semantics cannot be relied on to do any heavy lifting in metaphysics (given the minimal dose of realism which has been assumed herein), but—given the deep evident constitutive connections among meanings, concepts, beliefs—we cannot get very far in epistemology without engaging with semantics. Here again we return to a point from which we began in the Preface, for which this whole book aspires to serve as testament—namely,

that twentieth-century philosophy of language is the site of so many significant steps in the maturation of this ancient discipline. Work by Wittgenstein, Carnap, Quine, Kripke, and Kaplan has been marshalled into an argument that the constitutive *a priori* provides a strong compelling stance on Plato's ancient and epochal problem.

NOTES

1. Again, many would also hold that certain moral principles (e.g., one ought to keep one's promises) are also at once knowable *a priori*, analytically true, and metaphysically necessary. However, it is fair to say that these latter claims are more controversial than [1] or [2], since they presuppose a degree of metaphysical realism about social and moral phenomena that many philosophers reject.

2. This was extensively discussed in §3.2. For example, Kant's talk of 'containment' in defining analyticity (i.e., a judgment is analytic if the predicate is contained in the subject) seems to presuppose a quite outdated and untenable picture of concepts.

3. It has been objected to me that radiative decay, or events in the quantum void, afford counterexamples to [1], but I am not so sure that they should be conceded.

4. This principle says that, for certain privileged propositions P, if they are true, then they must be necessarily true. As discussed above in §4.3—see especially the quote from Kripke (1972: 159) reproduced there—Kripke argues that this principle holds for cases where essences are scientifically discovered (e.g., 'Heat is the motion of molecules', 'Gold is the element with atomic number 79'). There will be further discussion of such cases when we get to rows 2 & 4 in §7.3 below.

Could [1] be true but contingently so? Whether the laws of nature are in some sense contingent is one of the bones of contention between those who identify physical with metaphysical necessity and those who hold that physical necessities are a proper subset of the metaphysical necessities (cf. §3.1). Not to presuppose anything contentious on that front; the essential point here is just that no metaphysical conclusions follow from the point that there is no contradiction between the concepts of 'event' and 'uncaused'. (The fact that we might be able to conceive of the laws of nature being otherwise than they actually are does not yet show that other non-actual laws are metaphysically possible.)

5. Cf., e.g., Soames (2011), Casullo (2012).

6. For example, Kripke argues that the following are also necessary *a posteriori*:

4. Hesperus is Phosphorus.

5. Water is H₂O.

However, as first came up in chapter 2, the proper semantic analysis of the propositions expressed is quite controversial, and so I am steering clear of such cases.

7. That said, I acknowledge that there are many fine semantic distinctions available on the market, which would yield several distinct senses of the term 'analytic'

(cf., e.g., Russell [2008]). Perhaps on some such precisification, there may be a sense in which at least some of [1]–[3] should be understood as [N+, A+, AP-]. A direct-reference sort of semantic externalist (e.g., Salmon [1986], Soames [2002]) might even be inclined to try to classify [4] or [5], from note 6, as [N+, A+, AP-]. Hence, row 2 may have some defenders. I take it, though, that even if so, for the reasons given, then they would be a small unorthodox minority. (Some of these possibilities are explored in depth, from a different direction, in Sullivan [2010].)

8. Cf. §2.3. Kaplan (1989) and Stalnaker (2001) are among the seminal sources here; cf. Chalmers (2006) for an overview. Kripke's (1972) provocative discussions of 'contingent *a priori*' truth are also deeply relevant here; though, akin to note 6, I will avoid such examples as 'meter' and 'Neptune', as (i) in terms of the big picture, they do not add any challenge that is not already raised by the indexical cases, and (ii) I cannot do these examples any justice without getting into controversial theses within the philosophy of language that are not crucially relevant to the big modal picture. (For discussion of such cases see Soames [2003, Vol.II: Ch. 16].)

9. Other candidates for row 5 include examples involving the 'actuality' operator, such as 'If P is true, then it is actually true'. One classic source of the indexical approach to actuality is Lewis (1970); several variants have been developed.

10. See §3.3, and also Kripke (1972, 41–53; 1980, 15–20), for elaboration.

11. For a variety of arguments for the intelligibility of analyticity, cf. Grice & Strawson (1956), Fine (1994), Boghossian (1997), Katz (1997), Jackson (1998), Sober & Hylton (2000), Gertler (2002), Russell (2008), and Sullivan (2008).

12. Recall from §3.1 that there are some distinctive sorts of relativity, or context-shiftiness, to metaphysical modal claims. The key point is that it is rather distinct from the kind of framework-relativity that both the semantic and epistemic cases instance.

Bibliography

- Angeloni, R. (2012) "The Problem of 'Meaning Change' in Friedman's Notion of a Constitutive a Priori Principle." *Kairos* (5): 57–76.
- Armstrong, D. (1996) *A World of States of Affairs*. Cambridge University Press.
- Aune, B. (2008) *An Empiricist Theory of Knowledge*. Booksurge Publishing.
- Ayer, A.J. (1936) *Language, Truth, and Logic*. Dover Publications.
- Bach, K. (1987) *Thought and Reference*. Clarendon Press.
- . (2005) "Context *ex Machina*." In *Semantics vs. Pragmatics*, Z. Szabo, ed., Oxford University Press, pp. 15–44.
- Baehr, J. (2003) "A Priori and A Posteriori." In *Internet Encyclopedia of Philosophy*, J. Fieser, ed. (<http://www.iep.utm.edu/apriori>)
- Bealer, G. (1996) "A Priori Knowledge and the Scope of Philosophy." *Philosophical Studies* (81): 121–42.
- . (2000) "A Theory of the A Priori." *Pacific Philosophy Quarterly* (81): 1–30.
- Beall, J.C., and Restall, G. (2006) *Logical Pluralism*. Oxford University Press.
- Boghossian, P. (1997) "Analyticity." In Hale & Wright, ed., pp. 331–68.
- . (2000) "Knowledge of Logic." In Boghossian & Peacocke, ed., pp. 229–54.
- . (2010) Review of G. Russell (2008). *Australasian Journal of Philosophy* (89): 370–4.
- . (2011) "Williamson on the A Priori and the Analytic." *Philosophy and Phenomenological Research* (82): 488–97.
- Boghossian, P., and Peacocke, C., ed. (2000) *New Essays on the A Priori*. Oxford University Press.
- BonJour, L. (1998) *In Defense of Pure Reason*. Cambridge University Press.
- Burge, T. (1979) "Individualism and the Mental." *Midwest Studies in Philosophy* (4): 73–121.
- . (1986) "Intellectual Norms and the Foundations of Mind." *Journal of Philosophy* (83): 697–720.
- . (2000) "Frege on Apriority." In Boghossian & Peacocke, ed., pp. 11–42.

- . (2007) “Postscript to Individualism and the Mental.” In *Foundations of Mind*, Oxford University Press, pp. 151–81.
- Cappelen, H., and Lepore, E. (2005) *Insensitive Semantics*. Blackwell Press.
- Carnap, R. (1937) *Logical Syntax of Language*. Routledge & Keegan Paul.
- . (1947) *Meaning and Necessity*. University of Chicago Press.
- . (1950) “Empiricism, Semantics, Ontology.” *Revue Intenational de Philosophie* (4): 20–40.
- . (1963) “Replies and Systematic Expositions.” In *The Philosophy of Rudolph Carnap*, P. Schilpp, ed., Open Court Press, pp. 859–1013.
- Cartwright, N. (1983) *How the Laws of Physics Lie*. Oxford University Press.
- Cassam, Q. (2000) “Rationalism, Empiricism, and the A Priori.” In Boghossian & Peacocke, ed., pp. 43–64.
- Cassirer, E. (1919) *The Problem of Knowledge*. Yale University Press.
- Casullo, A. (2003) *A Priori Justification*. Oxford University Press.
- . (2012) “Analyticity, Apriority, Modality.” In *Continuum Companion to the Philosophy of Language*, M. Garcia-Carpintero and M. Kölbel, ed., Continuum International, pp. 228–50.
- Chalmers, D. (1996) *The Conscious Mind*. Oxford University Press.
- . (2006) “Two-Dimensional Semantics.” In *Oxford Handbook of the Philosophy of Language*, E. Lepore and B. Smith, ed., Oxford University Press, pp. 574–606.
- Chomsky, N. (1967) *Cartesian Linguistics*. Harper & Rowe.
- . (1993) *Language and Thought*. Moyer Bell.
- Chudnoff, E. (2011) “The Nature of Intuitive Justification.” *Philosophical Studies* (153): 313–33.
- . (2014) “The Rational Roles of Intuition.” In *Intuitions*, A. Booth and D. Rowbottom, ed., Oxford University Press, pp. 9–35.
- Coffa, A. (1991) *The Semantic Tradition from Kant to Carnap*. Cambridge University Press.
- Descartes, R. (1641) *Meditations on First Philosophy*. Hackett Publishing, 1979.
- Devitt, M. (2005) “There is no A Priori.” In *Contemporary Debates in Epistemology*, E. Sosa and M. Steup, ed., Blackwell Publishers, pp. 105–14.
- . (2011) “No Place for the A Priori.” In Shaffer & Veber, ed., pp. 9–32.
- DiPierri, G. (1992) “The Constitutive A Priori.” In Hanson & Hunter, ed., pp. 179–214.
- DiSalle, R. (2002) “Conventionalism and Modern Physics.” *Nous* (36): 169–200.
- Donnellan, K. (1970) “Proper Names and Identifying Descriptions.” *Synthese* (21): 335–58.
- . (1983) “Kripke and Putnam on Natural Kind Terms.” In *Knowledge and Mind*, C. Ginet and S. Shoemaker, ed., Oxford University Press, pp. 84–104.
- Dummett, M. (1978) *Truth and Other Enigmas*. Harvard University Press.
- Evans, G. (1982) *The Varieties of Reference*. Oxford University Press.
- Everett, J. (2015) “The Constitutive A Priori and the Distinction between Physical and Mathematical Possibility.” *Studies in the History and Philosophy of Science* (52): 139–52.

- Field, H. (2000) "A Priority as an Evaluative Notion." In Boghossian & Peacocke., ed., pp. 117–49.
- . (2005) "Recent Debates about the A Priori." In *Oxford Studies in Epistemology*, T. Gendler and J. Hawthorne, ed., Clarendon Press, pp. 69–88.
- Fine, K. (1994) "Essence and Modality." *Philosophical Perspectives* (8): 1–16.
- Flanagan, B. (2013). Analyticity and the Deviant Logician. *Acta Analytica* (28): 345–352.
- Franco, P. (2011) "Constitutivity and the Synthetic Method in Kant." *Hopos* (2): 321–53.
- Frege, G. (1884) *Foundations of Arithmetic*. J.L. Austin, transl. Blackwell, 1950.
- . (1892a) "On Sense and Reference." In *Translations from the Philosophical Writings of Gottlob Frege*, P. Geach and M. Black, transl., Blackwell Press, pp. 56–78.
- . (1892b) "On Concept and Object." In *Translations from the Philosophical Writings of Gottlob Frege*, P. Geach and M. Black, transl., Blackwell Press, pp. 42–55.
- Friedman, M. (1992) *Kant and the Exact Sciences*. Harvard University Press.
- . (2000) "Transcendental Philosophy and A Priori Knowledge." In Boghossian & Peacocke, ed., pp. 367–83.
- . (2001) *Dynamics of Reason*. CSLI Publications.
- . (2007) "The Evolution of the A Priori in Logical Empiricism." In *Cambridge Companion to Logical Empiricism*, A. Richardson and T. Uebel, ed., Cambridge University Press, pp. 91–116.
- . (2008) "Einstein, Kant, and the A Priori." *Royal Institute of Philosophy Supplement* (63): 95–112.
- . (2011) "Extending the Dynamics of Reason." *Erkenntnis* (75): 431–44.
- Gertler, B. (2002) "Explanatory Reduction, Conceptual Analysis, and Conceivability Arguments." *Nous* (36): 22–49.
- . (2012) "Understanding the Internalism-Externalism Debate." *Philosophical Perspectives* (26): 51–75.
- Gettier, E. (1963) "Is Justified True Belief Knowledge?" *Analysis* (23): 121–23.
- Gödel, K. (1944) "Russell's Mathematical Logic." In *The Philosophy of Bertrand Russell*, P. Schilpp, ed., Open Court Press, pp. 123–53.
- . (1947) "What is Cantor's Continuum Problem?" *American Mathematical Monthly* (54): 515–25.
- Goodman, N. (1952) "On Likeness of Meaning." *Analysis* (10): 1–7.
- Grayling, A.C. (1997) *Philosophical Logic* (3rd edition). Blackwell Press.
- Grice, H.P. (1987) "Retrospective Epilogue." In *Studies in the Ways of Words*, Harvard University Press.
- Grice, H.P., and Strawson, P. (1956) "In Defense of a Dogma." *Philosophical Review* (65): 141–58.
- Haack, S. (1974) *Deviant Logic*. Cambridge University Press.
- Hale, B. and Wright, C. ed. (1997) *Blackwell Companion to the Philosophy of Language*. Blackwell.

- Hanson, P., and Hunter, B., ed. (1992) *Return of the A Priori*. University of Calgary Press.
- Hawthorne, J. (2012) "Knowledge and Epistemic Necessity." *Philosophical Studies* (158): 493–501.
- Heinzman, G., and Stump, D. (2013) "Henri Poincare." *Stanford Encyclopedia of Philosophy*. (<https://plato.stanford.edu/entries/poincare/>)
- Horwich, P. (2000) "Stipulation, Meaning, and Apriority." In Boghossian & Peacocke, pp. 150–69.
- . (2006) *Reflections on Meaning*. Oxford University Press.
- Hume, D. (1748) *A Treatise of Human Nature*. Penguin Books, 1969.
- Jackson, F. (1998) *A Defense of Conceptual Analysis*. Oxford University Press.
- Kallestrup, J. (2012) *Semantic Externalism*. Routledge.
- Kant, I. (1781) *Critique of Pure Reason*. N. Kemp Smith, trnsl. MacMillan, 1929.
- . (1787) "Preface" to the 2nd edition of (1781).
- Kaplan, D. (1975) "Dthat." *Syntax & Semantics* (9): 221–53.
- . (1986) "Opacity." In *The Philosophy of W.V. Quine*. P. Schilpp, ed., Open Court Press, pp. 229–89.
- . (1989) *Themes from Kaplan*. Oxford University Press.
- Katz, J. (1997) "Analyticity, Necessity, and the Epistemology of Semantics." *Philosophy and Phenomenological Research* (57): 1–28.
- Kitcher, P. (2000) "A Priori Knowledge Revisited." In Boghossian & Peacocke, ed., pp. 65–91.
- Korkut, B. (2011) "The Real Virtue of Friedman's Neo-Kantian Philosophy of Science." *Philosophy of Science* (78): 1–15.
- Kripke, S. (1971) "Identity and Necessity." In *Identity and Individuation*, M. Munitz, ed., NYU Press, pp. 135–64.
- . (1972) *Naming and Necessity*. Harvard University Press.
- . (1979) "A Puzzle about Belief." In *Meaning and Use*, A. Margalit, ed., Reidel, pp. 239–83.
- . (1980) "Preface" to the 2nd edition of (1972).
- . (1986) "A Problem in the Theory of Reference." *Proceedings of the XVIIth World Congress of Philosophy*. Editions de Beffroi, pp. 241–47.
- Kuhn, T. (1964) *The Structure of Scientific Revolutions*. University of Chicago Press.
- Leibniz, G. (1704) *New Essay on Human Understanding*. Cambridge University Press, 1996.
- Lewis, C.I. (1923) "A Pragmatic Conception of the A Priori." *Journal of Philosophy* (20): 169–77.
- Lewis, D. (1970) "Anselm on Actuality." *Nous* (4): 175–88.
- . (1975) "Language and Languages." In *Minnesota Studies in the Philosophy of Science*, K. Gunderson, ed., University of Minnesota Press, pp. 3–35.
- . (1986) *On the Plurality of Worlds*. Blackwell Press.
- . (1998) Review of Armstrong (1996). *Times Literary Supplement* (4950): 30–33.
- Loar, B. (1991) "Can We Explain Intentionality?" In *Fodor and His Critics*, B. Loewer and G. Rey, ed., Blackwell Press, pp. 209–30.

- Locke, J. (1690) *An Essay Concerning Human Understanding*. Clarendon Press, 1984.
- Loewer, B. (1997) "A Guide to Naturalizing Semantics." In Hale & Wright, ed., pp. 108–26.
- Maddy, P. (2000) "Naturalism and the A Priori." In Boghossian & Peacocke, pp. 92–116.
- . (2007) *Second Philosophy*. Oxford University Press.
- Marcus, R. (1991) "A Backward Look at Quine's Animadversions on Modalities." In *Modalities*, Oxford University Press, pp. 215–32.
- Mates, B. (1950) "Synonymity." In *Semantics and the Philosophy of Language*, L. Linsky, ed., University of Illinois Press, pp. 111–38.
- Mill, J.S. (1843) *A System of Logic*. Longmans Press, 1947.
- Moore, G.E. (1903) *Principia Ethica*. Cambridge University Press.
- . (1925) "A Defense of Common Sense." In *Contemporary British Philosophy*, Muirhead, J. ed., Allen & Unwin, pp. 32–59.
- Munsat, S., ed. (1971) *The Analytic-Synthetic Distinction*. Wadsworth Publishing.
- Neale, S. (1990) *Descriptions*. MIT Press.
- . (2007) "Heavy Hands, Magic, and Scene Reading Traps." *European Journal of Analytic Philosophy* (3): 77–132.
- O'Grady, P. (2004) "Wittgenstein and Relativism." *International Journal of Philosophical Studies* (12): 315–37.
- Pap, A. (1946) *The A Priori in Physical Theory*. King's Crown Press.
- . (1958) *Semantics and Necessary Truth*. Yale University Press.
- Peacocke, C. (2000) "The Programme of Moderate Rationalism." In Boghossian & Peacocke, ed., pp. 255–85.
- . (2004) *The Realm of Reason*. Oxford: Oxford University Press.
- . (2006) "A Priori." In *Oxford Handbook of Contemporary Philosophy*, F. Jackson and M. Smith, ed., Oxford University Press, pp. 739–66.
- . (2011) "Understanding, Modality, Logical Operators." *Philosophy & Phenomenological Research* (82): 472–80.
- Pinker, S. (2011) "The Cognitive Revolution." *Harvard Gazette*. October 12.
- Plato (1928a) *The Works of Plato*. Translated by B. Jowett, edited by I. Edman. The Modern Library: 1956.
- . (1928b) "Seventh Letter." J. Harward, transl. *Classical Quarterly* (22): 143–54.
- Poincare, H. (1899) "On the Foundations of Geometry." *Revue de Métaphysique et de Morale* (7): 251–79.
- . (1900) "On the Principles of Geometry." *Revue de Métaphysique et de Morale* (8): 73–86.
- Prior, A. (1960) "The Roundabout Inference Ticket." *Analysis* (21): 38–9.
- Putnam, H. (1962) "It Ain't Necessarily So." *Journal of Philosophy* (59): 658–71.
- . (1975) "The Meaning of 'Meaning'." In *Minnesota Studies in the Philosophy of Science*, K. Gunderson, ed., University of Minnesota Press, pp. 225–71.
- . (1992) *Renewing Philosophy*. Harvard University Press.

- Quine, W.V. (1936) "Truth by Convention." In *Philosophical Essays for A.N. Whitehead*, O.H. Lee, ed., Longmans, pp. 70–99.
- . (1951) "Two Dogmas of Empiricism." *Philosophical Review* (60): 20–43.
- . (1956) "Quantifiers and Propositional Attitudes." *Journal of Philosophy* (53): 177–87.
- . (1960) *Word and Object*. MIT Press.
- . (1969) "Epistemology Naturalized." In "*Ontological Relativity*" and *Other Essays*, Columbia University Press, pp. 69–90.
- . (1970) *Philosophy of Logic*. Prentice Hall.
- . (1991) "Two Dogmas Revisited." *Canadian Journal of Philosophy* (21): 265–74.
- . (1992) *Pursuit of Truth*. Harvard University Press.
- Quine, W.V., and Ullian, J. (1970) *Web of Belief*. Random House.
- Quinton, A. (1963) "The A Priori and the Analytic." *Proceedings of the Aristotelian Society* (64): 31–54.
- Railton, P. (2000) "A Priori Rules." In Boghossian & Peacocke, pp. 170–96.
- . (2003) *Facts, Values, and Norms*. Cambridge University Press.
- Recanati, F. (1993) *Direct Reference*. Blackwell Press.
- . (2004) *Literal Meaning*. Cambridge University Press.
- Reichenbach, H. (1920) *The Theory of Relativity and A Priori Knowledge*. Springer.
- . (1953) *Nomological Statements and Admissible Operations*. North Holland.
- Richardson, A. (2002) "Narrating the History of Reason Itself: Friedman, Kuhn, and a Constitutive a Priori for the Twenty-First Century." *Perspectives on Science* (10): 253–74.
- Roca-Royes, S. (2011) "Essential Properties and Individual Essences." *Philosophy Compass* (6): 65–77.
- Russell, B. (1912) *The Problems of Philosophy*. Allen & Unwin.
- . (1918) "The Philosophy of Logical Atomism." *Monist* (28/29).
- Russell, G. (2008) *Truth in Virtue of Meaning*. Oxford University Press.
- . (2012) "Analyticity for Externalist Languages." In *New Waves in the Philosophy of Language*, S. Swoyer, ed., Palgrave MacMillan, pp. 186–205.
- Salmon, N. (1986) *Frege's Puzzle*. MIT Press.
- . (1989) "Illogical Belief." *Philosophical Perspectives* (3): 243–85.
- Schaffer, M. (2011) "The Constitutive A Priori and Epistemic Justification." In Schaffer & Veber, pp. 193–210.
- Schaffer, M., and Veber, M., ed. (2011) *What Place for the A Priori?* Open Court Press.
- Schiffer, S. (2003) *The Things We Mean*. Oxford University Press.
- Searle, J. (1969) *Speech Acts*. Cambridge University Press.
- . (1978) "Literal Meaning." *Erkenntnis* (13): 207–24.
- Sober, E., and Hylton, P. (2000) "Quine's Two Dogmas." *Proceedings of the Aristotelian Society* (74): 237–80.
- Soames, S. (2002) *Beyond Rigidity*. Oxford University Press.
- . (2003) *Philosophical Analysis in the Twentieth Century*. Princeton University Press.

- . (2011) “Kripke on Epistemic and Metaphysical Possibility.” In *Saul Kripke*, A. Berger, ed., Cambridge University Press, pp. 78–99.
- Stalnaker, R. (1972) “Assertion.” *Syntax & Semantics* (9): 315–32.
- . (1984) *Inquiry*. MIT Press.
- . (2001) *Content and Context*. Oxford University Press.
- Stanley, J. (2007) *Language in Context*. Oxford University Press.
- Stern, D. (2007) “Wittgenstein, the Vienna Circle, and Physicalism.” In *Cambridge Companion to Logical Empiricism*, A. Richardson and T. Uebel, ed., pp. 305–31.
- Strawson, P. (1959) *Individuals*. Methuen Press.
- Stump, D. (2003) “Defending Conventions as Functionally A Priori Knowledge.” *Philosophy of Science* (20): 1149–60.
- . (2011) “A Reconsideration of the Status of Newton’s Laws.” In Schaffer & Veber, ed., pp. 177–92.
- . (2015) *Conceptual Change and the Philosophy of Science*. Routledge Press.
- Sullivan, A. (2003a) “Shareability and Objectivity.” *Ratio* (16): 251–71.
- . (2003b) “‘Paging Dr Lauben!’ Some Questions about Individualism and Competence.” *Philosophical Studies* (115): 201–24.
- . (2008) “Truth in Virtue of Meaning.” *Canadian Journal of Philosophy* (38): 373–97.
- . (2010) “Millian Externalism.” In *New Essays on Singular Thought*, R. Jeshion, ed., Oxford University Press, pp. 246–69.
- . (2012) *Reference and Structure in the Philosophy of Language*. Routledge Press.
- . (2015a) “The Future of the Proposition.” *Language and Communication* (40): 14–23.
- . (2015b) “What do Deviant Logicians Show about the Epistemology of Logic?” *Acta Analytica* (30): 179–91.
- Sullivan, A., ed. (2003) *Logicism and the Philosophy of Language*. Broadview Press.
- Taylor, K. (2007) “A Little Sensitivity Goes a Long Way.” In *Content-Sensitivity and Semantic Minimalism*, G. Preyer and G. Peter, ed., Oxford University Press, pp. 63–93.
- Uebel, T. (2012) “De-Synthesizing the Relative A Priori.” *Studies in the History and Philosophy of Science* (43): 7–17.
- Weinberg, J. (2014) “The Prospects for Experimental Rationalism.” In *The A Priori in Philosophy*, A. Casullo and J. Thurow, ed., Oxford University Press, pp. 92–110.
- White, M. (1950) “The Analytic and the Synthetic: An Untenable Dualism.” In *Semantics and the Philosophy of Language*, L. Linsky, ed., University of Illinois Press, pp. 272–86.
- Wilkforss, A. (2010) Review of G. Russell (2008). *Notre Dame Philosophical Reviews*. (<http://ndpr.nd.edu/news/23867-truth-in-virtue-of-meaning/>)
- Williamson, T. (2006) “Conceptual Truth.” *Proceedings of the Aristotelian Society* (80): 1–41.
- . (2008) *The Philosophy of Philosophy*. Oxford: Blackwell Press.
- . (2009) Reply to Kornblith. *Analysis* (69): 123–26.

- . (2011) Reply to Boghossian. *Philosophy & Phenomenological Research* (82): 498–506.
- . (2014) “How Deep is the Distinction between A Priori and A Posteriori Knowledge?” In *The A Priori in Philosophy*, A. Casullo and J. Thurow, ed., Oxford University Press, pp. 291–312.
- Wittgenstein, L. (1921) *Tractatus Logico-Philosophicus*. Routledge & Keegan Paul.
- . (1953) *Philosophical Investigations*. Blackwell Press.
- . (1969) *On Certainty*. Blackwell Press.
- . (1979) *Notebooks 1914–16*. 2nd edition. University of Chicago Press.

Index

- a priori*, vii–x, 3–14, 16–23, 27–34, 40–44, 56–57, 68–69, 75–96, 99–108, 115–18, 120, 123–32, 132–39, 139–44, 153–55, 161–76;
- acquaintance–based vs.
understanding–based, 18–19, 40–44, 77, 84–85, 87, 91, 94, 103, 108–9, 115–16, 120, 129, 136, 139, 143, 153–55, 171;
- constitutive, vii, 6, 12, 19–22, 40–43, 63, 87, 94, 99, 103–8, 120, 123–32, 132–39, 139–44, 153–55, 161–74;
- contingent, viii, 5, 73, 90–91, 96, 105, 126, 175;
- synthetic, 12–14, 19, 23, 61, 90–92, 103, 164, 167–68, 173
- absolutism, 100–2, 104, 128, 141, 153
- analytic, vii–x, 4–6, 11–12, 20, 22, 30–33, 40–44, 55–63, 66–71, 72, 89–92, 94, 115, 120, 161–76
- Aristotle, 14, 26, 51, 99, 154
- Boghossian, Paul, x, 18, 30–32, 44, 66–70, 72–73, 93, 95, 101, 121, 133, 158, 176
- Burge, Tyler, 22, 23, 44, 69, 83, 86, 88, 96, 119, 121, 151
- Carnap, Rudolph, viii, 12, 20, 35, 101, 105–8, 123, 128–32, 136, 139, 141, 143, 157, 170, 175
- Casullo, Albert, x, 22, 87, 95, 101, 175
- Coffa, J.A., 12, 44, 59–60, 72–73, 101, 124, 129
- competence, 21, 38, 62, 102–4, 109–15, 151–53
- conceivability, 23, 28, 132, 172
- conceptual evolution, 20, 41–43, 61, 87–88, 95, 100, 108, 128, 130, 139–44, 166, 172–73
- convention, 23, 28–33, 61, 116–17, 122, 129
- deference, 21, 112–19, 122, 144–53, 157–58, 166, 172
- Descartes, Rene, vii, 4, 10, 18, 22, 75, 83–86, 95, 103, 167
- empiricism, 3, 10, 12, 17–19, 76–77, 82–83, 100–2, 116–17, 133, 155
- epistemology, viii, 3–14, 14–17, 23, 75–96, 100, 115, 141, 153–55, 161–75
- essence, 51–53, 96, 117, 147, 157, 165–66, 175
- extension, 21, 34–35, 39–40, 43, 63, 109–15, 116–17, 147, 172

- externalism, 20–21, 24, 40, 99, 109–15,
116–21, 123, 145–54, 162, 170,
172, 176
- Field, Hartery, 95, 101, 120, 143–44
- framework, 20, 22, 43, 79, 87–88,
101–4, 104–8, 115–20, 123, 125,
128–32, 137–38, 139–48, 153–55,
161–75, 176
- Frege, Gottlob, vii, 13, 21–22, 35, 38,
62–63, 66, 69, 72–73, 75, 77, 81,
86, 95–96, 100, 109, 114, 121,
149, 163
- Friedman, Michael, viii, 94, 101, 105–8,
120, 130–32, 156
- Hume, David, 4, 12–13, 17, 59, 65, 69,
73, 81, 117, 133, 154, 174
- immunity to counterexample, 3–6, 6–12,
20, 31–32, 36–37, 40–44, 55–57,
62, 71, 75–83, 86–93, 99–109,
114–18, 124–30, 132–41, 143–44,
153–55, 158, 161–75
- incomplete mastery, 21, 112–14,
144–53, 158
- indefeasibility, 22, 87, 93, 95, 127
- indexicality, vii, 5, 20–21, 33, 36–40,
44–45, 84, 167–70, 176
- innateness, 77–78, 86–88, 95
- intuition, 6, 9–10, 16–19, 84, 114, 117,
120, 147, 153–55, 171, 173–74;
rational, 9–10, 17–19, 23, 43–44,
55, 76–79, 81–85, 90, 93, 103,
155, 171, 173–74;
semantic, 16–19, 43–44, 55, 63, 68,
77, 79, 85, 94, 103–4, 109, 115,
117, 120, 172, 173
- justification, 3–4, 6, 17–20, 23, 43–44,
62, 75–83, 87, 94–96, 102–4,
107, 114–15, 118, 124, 129–30,
132–39, 143, 153–55, 158, 161–74
- Kant, Immanuel, vii–viii, 4, 12–14, 20,
22–23, 61–62, 66, 75, 77, 81, 83,
86, 92, 94, 96, 100–101, 104,
117, 123, 130–32, 154, 163–64,
174
- Kaplan, David, viii, x, 5, 23, 35, 37–39,
45, 73, 121, 167, 175
- knowledge, 3–4, 7–10, 15–16, 17–18,
43–44, 75–83, 87, 95, 96, 116–17,
153–55, 161–74
- Kripke, Saul, vii–x, 5, 17, 20, 21,
22–23, 30, 52, 71–73, 91–92,
96, 99, 109–15, 117, 121, 141,
149–54, 162–66, 170, 174–76
- Maddy, Penelope, 18, 93–94, 121,
143–44, 157
- meaning, 14–17, 21, 24, 28–33, 34–40,
43–45, 57–63, 66–71, 109–15,
121, 143, 147, 150, 155, 171
- metaphysics, viii, 4–10, 14–17, 23,
40–44, 49–55, 64–66, 115, 124,
141, 154, 161–74
- modality, 36–37, 40, 50–55
- natural kinds, 21, 39–42, 54, 96,
110–15, 116–20, 122, 144–53,
165–66, 170–71, 176
- naturalism, 10, 18, 23, 58, 93, 99, 106,
121
- necessary, vii–viii, 4–14, 28, 30–33,
40–44, 49–55, 64–66, 89–92,
96, 115, 117, 129–30, 132, 155,
161–75;
a posteriori viii, 5, 90–91, 96, 117,
155, 164–66;
physical vs. metaphysical necessity,
28, 53–54, 175
- Pap, Arthur, viii, 22, 49, 94, 101, 105–8,
128–30, 169
- Peacocke, Christopher, x, 19, 22, 75, 93,
95, 101, 133, 143
- Plato, vii, 4, 7–11, 19, 49, 75, 77, 80,
84, 95, 154, 161;
Plato's problem, 8–12, 19, 59, 73,
77, 101–4, 108, 115, 120, 128,
139, 153–55, 161–74

- Poincare, Henri, viii, 129–30, 133, 140, 156, 170
- proper names, 39–40, 96, 110–15, 149, 175
- proposition, 26–28, 35–40, 40–45, 56, 60, 72, 77–79, 124–25, 172
- Quine, W.V., vii–x, 4, 10, 17, 18, 20, 22–23, 26, 32, 66–70, 73, 75, 93, 99–102, 115, 119–20, 130, 134, 141–43, 154, 162, 175
- Railton, Peter, viii, 101, 105, 142, 157
- rationalism, 9–12, 18–20, 23, 76–77, 116–17, 133, 155
- reference, 21, 24, 40, 63, 110–15, 121, 144–53, 155;
determiner, 40, 63, 114, 121
- Reichenbach, Hans, viii, 101, 128–30, 156, 170
- relativism, 128, 130
- revisability, ix, 19–23, 40–44, 63, 71, 86–88, 99–109, 115–20, 123, 129, 137–57, 162, 172
- Russell, Bertrand, 10, 18, 38, 69, 81, 109, 114, 124–25, 149
- Russell, Gillian, x, 22, 63, 66, 73, 151, 176
- semantics, viii, 4–6, 11–12, 14–17, 23, 34–40, 40–44, 57–63, 100, 109–15, 141, 154, 163
- skepticism, vii, 3, 10–12, 18, 57–58, 63–71, 92–95, 100–5, 120, 124, 134–35, 153
- Stump, David, viii, 94, 101, 105–8, 120, 130–32, 156
- transparency, 114, 117, 119, 121, 146–48, 151–53, 157–58
- truth, 25–26, 44, 73, 77–79
- Williamson, Timothy, ix, 72, 121, 133–39, 151, 156, 158
- Wittgenstein, Ludwig, viii, 12, 20, 23, 69, 101, 105–13, 121, 123–31, 133, 136, 141–42, 155–57, 170, 175

About the Author

Arthur Sullivan is associate professor of philosophy at Memorial University of Newfoundland. He works primarily in the philosophy of language, but has also published research in related areas of the philosophy of mind, metaphysics, epistemology, and the history of philosophy. His papers have appeared in journals such as *Journal of Pragmatics*, *Philosophers' Imprint*, *Philosophical Studies*, and 'Synthese', and his previous books include an edited volume of work by Frege and Russell, *Logicism and the Philosophy of Language* (2003), and a research monograph in the theory of reference *Reference and Structure: A Defense of the Russellian Orthodoxy* (2013).

