

Constructing the World

David J. Chalmers



OXFORD
UNIVERSITY PRESS



Contents

<i>Acknowledgments</i>	xi
<i>Introduction</i>	xiii
<i>How to Read this Book</i>	xv
1 Scrutability and the <i>Aufbau</i>	1
1 Primitive concepts	1
2 Objections to the <i>Aufbau</i>	7
3 From Definitional Scrutability to A Priori Scrutability	12
4 From descriptions to intensions	16
5 The scrutability base	20
6 Scrutability and the <i>Aufbau</i>	25
First Excursus: Scrutability and Knowability	29
Second Excursus: The Inscrutability of Reference and the Scrutability of Truth	34
2 Varieties of Scrutability	39
1 Scrutability theses	39
2 Sentences or propositions?	42
3 Inferential scrutability	47
4 Conditional scrutability	53
5 A priori scrutability	58
6 Generalized scrutability	60
7 Idealization	62
8 Objections from idealization	64
Third Excursus: Sentential and Propositional Scrutability	72
Fourth Excursus: Warrants and Support Structures	92
Fifth Excursus: Insulated Idealization and the Problem of Self-Doubt	101
3 Adventures with a Cosmoscope	108
1 A scrutability base	108
2 The Cosmoscope argument	114
3 The argument from elimination	120
4 The argument from knowability	125
5 Inferential scrutability with a Cosmoscope	134
6 Conditional scrutability	138
7 The objection from recognitional capacities	139
8 The objection from counterfactuals	148
Sixth Excursus: Totality Truths and Indexical Truths	151



4	The Case for A Priori Scrutability	157
1	From Conditional to A Priori Scrutability	157
2	The argument from suspension of belief	159
3	The argument from frontloading	160
4	Causal roles, mediating roles, and justifying roles	167
5	Generalized A Priori Scrutability	169
6	Objections from self-knowledge	171
7	Objections from theories of concepts and reference	173
8	Objections from acquaintance and from nonpropositional evidence	176
9	The objection from empirical inference	181
	Seventh Excursus: Varieties of Apriority	185
	Eighth Excursus: Recent Challenges to the A Priori	193
5	Revisability and Conceptual Change	199
1	Introduction	199
2	The arguments of “Two Dogmas”	200
3	Carnap on intensions	204
4	A Carnapian response	205
5	Refining Carnap’s account	207
6	A Bayesian analysis of holding-true	211
7	A Bayesian analysis of revisability	214
8	Quinean objections	217
9	Conclusion	224
	Ninth Excursus: Scrutability and Conceptual Dynamics	226
	Tenth Excursus: Constructing Epistemic Space	233
	Eleventh Excursus: Constructing Fregean Senses	244
6	Hard Cases	259
1	Introduction	259
2	Mathematical truths	261
3	Normative and evaluative truths	264
4	Ontological truths	267
5	Other philosophical truths	271
6	Modal truths	273
7	Intentional truths	274
8	Social truths	279
9	Deferential terms	280
10	Names	282
11	Metalinguistic truths	283
12	Indexicals and demonstratives	285
13	Vagueness	288
14	Secondary qualities	289
15	Macrophysical truths	290
16	Counterfactual truths	298
17	Conclusion	299
	Twelfth Excursus: Scrutability and the Unity of Science	301



CONTENTS

ix

7	Minimizing the Base	312
1	Introduction	312
2	Heuristics	313
3	Microphysical expressions	319
4	Color, other secondary qualities, and mass	321
5	Spatiotemporal expressions	325
6	Causal and nomic expressions	336
7	Phenomenal expressions	340
8	Compression using laws	344
9	Quiddities	347
10	Other expressions	353
11	Packages	357
	Thirteenth Excursus: From the <i>Aufbau</i> to the Canberra Plan	362
	Fourteenth Excursus: Epistemic Rigidity and Super-Rigidity	366
8	The Structure of the World	379
1	Principled scrutability bases	379
2	Definitional Scrutability (and conceptual analysis)	381
3	Analytic and Primitive Scrutability (and primitive concepts)	385
4	Narrow Scrutability (and narrow content)	393
5	Acquaintance Scrutability (and Russellian acquaintance)	399
6	Fundamental Scrutability (and the mind–body problem)	404
7	Structural Scrutability (and structural realism)	409
8	Generalized Scrutability (and Fregean content)	423
	Summation: Whither the <i>Aufbau</i> ?	426
	Fifteenth Excursus: The Structuralist Response to Skepticism	431
	Sixteenth Excursus: Scrutability, Supervenience, and Grounding	441
	Seventeenth Excursus: Explaining Scrutability	461
	<i>Glossary</i>	468
	<i>Bibliography</i>	475
	<i>Index</i>	485



Introduction

In 1814, Pierre-Simon Laplace wrote:

An intellect which at a certain moment would know all forces that set nature in motion, and all positions of all items of which nature is composed, if this intellect were also vast enough to submit these data to analysis, it would embrace in a single formula the movements of the greatest bodies of the universe and those of the tiniest atom; for such an intellect nothing would be uncertain and the future just like the past would be present before its eyes.

Laplace suggests that given the right basic information, and sufficiently powerful reasoning, all truths about the universe can be determined. For Laplace, this basic information included truths about the fundamental laws of physics and truths about the location of all fundamental entities at a time. Let us call these the *Laplacean truths*. The reasoning requires an idealized “vast enough intellect”, which we might call a *Laplacean intellect*. A Laplacean intellect who knows all the Laplacean truths is a *Laplacean demon*.

The key claim in Laplace’s text is that a Laplacean demon would be uncertain of nothing. In effect, Laplace is saying that for any proposition that the demon can entertain, the demon will not be uncertain about that proposition. Or in a small variation: for any proposition the demon can entertain, the demon will be in a position to know whether it is true.

Suppose that there will be an election tomorrow. I can entertain the proposition that the left-wing candidate will win, the proposition that the right-wing candidate will win, and the proposition that the third-party candidate will win. If Laplace’s thesis is right, a Laplacean demon in my shoes will be able to know which if any of these propositions is true. If the left-wing candidate will win, the demon will be in a position to know it; if the right-wing candidate will win, the demon will be in a position to know it; if the third-party candidate will win, the demon will be in a position to know it.

Laplace’s thesis is an instance of what I call a *scrutability* thesis. It says that the world is in a certain sense comprehensible, at least given a certain class of basic truths about the world. In particular, it says that all truths about the world are *scrutable* from some basic truths. This means roughly that there is a connection in the realm of knowledge between the basic truths and all the rest: given the basic truths, the rest of the truths can be determined.

We might then put a version of Laplace’s thesis as follows:

Laplacean Scrutability: For all true propositions p , a Laplacean intellect who knew all the Laplacean truths would be in a position to know that p .

In the years since Laplace wrote, Laplace's demon has come in for something of a battering. But I think that there remains much value in Laplace's pregnant idea. One can extract some of the value by examining the problems that arise for Laplace's thesis, and by reformulating the scrutability thesis in a way that avoids them.

One sort of problem arises from the information that Laplace allows in the base. Most famously, the failure of determinism in quantum mechanics suggests that the demon could not predict the future just from facts about physical laws and about the present. It may be, for example, that futures in which the left-wing candidate, the right-wing candidate, and the third-party candidate win are all left open by these facts. All three futures could evolve from the present state of the world given the right sort of quantum-mechanical evolution.

There are other limitations. Many have argued that complete physical information is not enough to know all truths about the mind: if Laplace's demon has never experienced colors, for example, it will not know what it is like to see red. Others have argued that complete objective information is not enough to determine perspectival truths about the current time, or one's own identity: even given complete physical information, Laplace's demon might not know that today is Tuesday. Others find gaps for mathematics, morality, and other areas.

To avoid these problems, however, we need only give Laplace's demon more information than Laplace allows. To accommodate nondeterminism, we might give the demon full information about the distribution of fundamental physical entities throughout space and time. To handle problems involving the mind and the self, we might give the demon information about consciousness or the principles governing it, along with information about its own location in spacetime. If there are gaps for mathematics or morality, we can give Laplace's demon mathematical and moral principles as well. It is not clear precisely what information is required, but here the key claim is that there is *some* limited class of base truths that will allow Laplace's demon to do its work.

We might say that a *compact* class of truths is a set of truths that involves only a limited class of concepts and that avoids trivializing mechanisms such as coding the entire state of the world into a single number. I will elaborate on this rough characterization in the first chapter. For now, we can say that the class of physical truths will be a compact class, as will the expanded class of truths suggested above. We might then put a generalized Laplacean thesis as follows:

Inferential Scrutability: There is a compact class of truths such that for all true propositions p , a Laplacean intellect who knew all the truths in that class would be in a position to know that p .

Inferential Scrutability allows a broader class of base truths than Laplacean Scrutability, but otherwise it shares a similar form. In both theses, the idea is that if the demon knew all the basic truths, it could come to know all the truths, perhaps by inference from those basic truths. For the demon to know all the basic truths, they must be true in the demon's own world. So Inferential Scrutability in effect requires that the demon inhabits our world or one very much like it, knows all the basic truths about it, and comes to know all truths from there.

This requirement gives rise to a second sort of problem for Laplace's demon. In the actual world, we may suppose, one truth is that there are no Laplacean demons. But no Laplacean demon could know that there are no Laplacean demons. To avoid this problem, we could require the demon to know all truths about its own modified world rather than about the actual world. But now the demon has to know about itself, and a number of paradoxes threaten. There are paradoxes of complexity: to know the whole universe, the demon's mind needs to be as complex as the whole universe, even though it is just one part of the universe. There are paradoxes of prediction: the demon will be able to predict its own actions and then try to act contrary to the prediction. And there are paradoxes of knowability: if there is any truth q that the demon never comes to know, perhaps because it never entertains q , then it seems that the demon could never know the true proposition that q is a truth that it does not know.

To avoid these paradoxes, we can think of the demon as lying outside the world it is trying to know. Or better, we can think of the demon as contemplating the universe conditionally: *if* the Laplacean truths obtain, *then* this is what follows. Even if our own world does not contain a demon, we can still ask what a demon in some other world could come to know about our world, if it were given the relevant information in conditional form. Such a demon need not contemplate its own existence. What results is a conditional version of the scrutability thesis.

Conditional Scrutability: There is a compact class of truths such that for all true propositions p , a Laplacean intellect would be in a position to know that *if* the truths in that class obtain, then p .

We can make one final change. A key element of Laplace's idea is that the Laplacean truths are *all* the truths that the demon needs. No other empirical information is needed for the demon to do its job. Here, the idea is that to know the conditional above—*if* the basic truths hold, then p holds—the demon does not need any further empirical information in the background. That is, to know the conditional, the demon need not rely on a posteriori sources such as perception, introspection, or testimony. The demon can know the conditional a priori, or with justification independent of experience. We might put this as follows.

A Priori Scrutability: There is a compact class of truths such that for all true propositions p , a Laplacean intellect would be in a position to know a priori that if the truths in that class obtain, then p .

The three preceding theses are all scrutability theses. They say that there is a compact class of basic truths from which all truths can be determined, given sufficiently powerful reasoning. Versions of the A Priori Scrutability thesis will be my central focus in this book, but the other theses above will also play a role.

All sorts of questions immediately arise. How can the scrutability thesis be made precise? Why should we believe it, and how can one argue for it? Just which truths are among the basic truths, and how small can the basis be? What about hard cases, such as knowledge of social truths, or moral truths, or mathematical truths? What does the scrutability thesis tell us about language, thought, knowledge, and reality? All of these questions are pursued in this book.

I suspect that to many readers, the scrutability theses just discussed will seem obvious. I hope that they at least seem plausible to many more. But to paraphrase Russell, philosophy is the art of moving from obvious premises to interesting conclusions. Even if scrutability theses are obvious, there are many interesting conclusions downstream from them. Of course theses requiring only a compact base do not do everything that Laplace's stronger thesis could do. If a demon is given a full specification of how basic physical entities are distributed throughout space and time, for example, then its ability to predict the future is less impressive than it would have been for Laplace's original demon. I think that nevertheless, the thesis is powerful, because of its many applications.

Applications to epistemology, the study of knowledge, are perhaps the most obvious. For example, the scrutability thesis is at least a cousin of the knowability thesis, the thesis that all truths can be known. In addition, I will argue later that a version of the scrutability thesis can help with the problem of skepticism about the external world.

There are also applications in many other areas. In metaphysics, specific versions of the scrutability thesis can be used to help determine what is true and what is fundamental. In the philosophy of science, the scrutability thesis can be used to shed light on reductive explanation and the unity of science. In the philosophy of mind, the scrutability thesis can be used to help understand primitive concepts and the content of thought. And perhaps most importantly, the thesis has powerful applications in the philosophy of language, helping us to analyze notions of meaning and content that are tied to thought and knowledge.

In fact, the scrutability framework bears directly on many of the central debates in philosophy. One version of the thesis can be used to defend a Fregean approach to meaning (an analysis of meaning grounded in rationality and the a

priori) over a purely Russellian approach (an analysis grounded in reference and the external world). Another can be used to defend internalism about mental content, defining a sort of content that is largely intrinsic to the subject, against a strong externalism on which all content depends on the environment. Another can be used as a key premise in an argument against materialism about consciousness. Another can be used to deflate many traditional skeptical arguments about knowledge. Another can be used to support a version of structural realism about science.

Different versions of the scrutability thesis are relevant to different applications, so the issues do not all stand and fall together.¹ But in different ways, scrutability provides a powerful fulcrum through which we can gain leverage on these issues. In some cases, one can make related arguments without a direct appeal to scrutability, so the conclusions are not wholly beholden to the scrutability framework and can be cast in different terms. But in every case, thinking in terms of scrutability reframes the issues in a way that can make old, murky problems a little clearer.

The scrutability framework tends in a direction contrary to a number of trends in post-1950 philosophy: trends including direct reference theories of meaning, externalism about mental content, and rejection of the analytic/synthetic distinction. In various respects, it helps to support ideas from an earlier era in philosophy. It supports Gottlob Frege's distinction between sense and reference, and helps provide a concrete account of what Fregean senses are. It coheres well with Bertrand Russell's ideas about constructions of the external world and about the role of acquaintance in thought and knowledge. And above all, it provides support for many key ideas of the great logical empiricist, Rudolf Carnap.

In many ways, Carnap is the hero of this book. Like the other twentieth-century logical empiricists, he is often dismissed as a proponent of a failed research program. But I am inclined to think that Carnap was fundamentally right more often than he was fundamentally wrong. I do not think that he was right about everything, but I think that many of his ideas have been underappreciated. So one might see this project, in part, as aiming for a sort of vindication.

The title of this book is a homage to Carnap's 1928 book *Der Logische Aufbau der Welt*, usually translated as either *The Logical Construction of the World* or *The Logical Structure of the World*. The title (like Carnap's?) should be heard as self-consciously absurd. I am not really constructing the world. But one can see the

¹ The application to Fregeanism requires a generalization of the A Priori Scrutability thesis already stated, while the other four applications respectively involve what I later call Narrow, Fundamental, Nomic, and Structural Scrutability respectively. The first application is outlined in the eleventh excursus, and the other four applications are discussed in and around chapter 8.

current book as trying to carry off a version of Carnap's project in the *Aufbau*: roughly, constructing a blueprint of the world, or at least constructing a blueprint for a blueprint, by providing a vocabulary in which such a blueprint can be given, and making a case that the blueprint would truly be a blueprint for the world. More specifically, the aim is to specify the structure of the world in the form of certain basic truths from which all truths can be derived. To do this, I think one has to expand Carnap's class of basic truths and change the derivation relation, just as we had to for Laplace. But with these changes made, I think that the project is viable and that some of the spirit of the *Aufbau* remains intact.

I did not set out to write a Carnapian book. Instead, the connections between my project and Carnap's crept up on me to the point where they could not be ignored. The connections to Carnap go beyond the *Aufbau*. The approach to Fregean sense in terms of intensions is very much a descendant of Carnap's approach in *Meaning and Necessity*. The reply to Quine in chapter 5 can be seen as an adaptation of Carnap's analysis in "Meaning and Synonymy in Natural Language". My approaches to the unity of science, to skepticism, to inferentialism, and to verbal disputes all have something in common with different elements of Carnap's work. In some cases I was not conscious of the connection to Carnap until well into the process, but his presence here is clear all the same.

I should not overstate the extent to which my views and my motivations are Carnap's. I am not a logical empiricist or a logical positivist. I do not share Carnap's sometime inclination toward verificationism and phenomenalism. Where Carnap invokes a semantic notion of analyticity, I invoke an idealized epistemological notion of apriority. Logic plays a less central role for me than for Carnap, and unlike him, I eschew explicit definitional constructions. Carnap would not have approved of my views on the mind–body problem. Where Carnap saw the *Aufbau* as an attempt to make the content of science wholly objective and communicable, vindicating science serves less as a motivation for me, and my version of the project has subjective and nonstructural elements right in the base.

So this book picks up only on certain strands in Carnap, and not on his project as a whole. To oversimplify, one might say that where Carnap leans toward empiricism, I lean toward rationalism. The project as a whole might be seen as a sort of Carnapian rationalism. To some, that label might seem oxymoronic, but this just brings out that there is more to Carnap than traditional caricatures may suggest.

That said, I would like to think that those who share more of Carnap's empiricism than I do will find that there are still many elements of the current picture that they can accept. Later in the book, I discuss ways in which a version of this project might be used to vindicate something quite close to the Carnapian picture, coming as close as possible to the structural and definitional picture in the *Aufbau*.

Here is roughly what happens in this book. Chapter 1 introduces the project using the *Aufbau* as a guide. I go over various objections to the *Aufbau*, and sketch a version of the project that has the potential to overcome all these objections. This chapter in effect motivates and gives an overview of the project of the book as a whole. Chapter 2 goes over preliminaries, formulating scrutability theses in detail and addressing a number of other preliminary issues.

Chapters 3 and 4 mount the core arguments for scrutability. I argue for a limited scrutability thesis concerning the scrutability of all “ordinary truths” from a certain base. Chapter 3 focuses on Inferential and Conditional Scrutability, using a hypothetical device, the Cosmoscope, to make things vivid. Chapter 4 extends these arguments to A Priori Scrutability. Many epistemological issues come up along the way in these chapters, and numerous objections are addressed.

Chapter 5 uses the framework to respond to Quine’s arguments against analyticity and apriority, by providing an analysis of conceptual change. Along the way, it develops a notion of meaning inspired by Carnap and grounded in the scrutability framework. Chapter 6 extends the arguments of chapters 3 and 4 to the scrutability of all truths, by considering various “hard cases” such as mathematical truths, normative truths, intentional truths, ontological truths, and many others.

Chapters 7 and 8 investigate the character of a scrutability base. Chapter 7 tries to whittle down the base to the smallest possible class, proceeding through various domains to see whether they involve primitive concepts and need to be in the base or whether they can be eliminated. Chapter 8 builds on this to investigate the prospects for certain principled scrutability theses, in part to see to what extent the projects of Carnap and Russell can be vindicated, and in part to develop various applications. I see these two chapters as perhaps the central chapters of the book. Chapter 7 goes over many important issues concerning what should be in the base, while chapter 8 gives a sense of the upshot and rewards of the project. A summation after chapter 8 reviews the prospects for *Aufbau*-like projects, arguing that projects in the spirit of Carnap and Russell look surprisingly good.

Along the way, a series of excursuses after each chapter explore all sorts of connected issues. “Excursus” is usefully ambiguous between “a detailed discussion of a particular point in a book, usually in an appendix” and “a digression in a written text”. Some of my excursuses are detailed discussions of points within the framework. Some of the more important excursuses in this group are the third (on sentential and propositional scrutability), the fifth (on insulated idealization), and the fourteenth (on epistemic rigidity and super-rigidity). Others are digressions that draw connections to the philosophical literature or develop applications. The more important excursuses in this group include the eleventh (on meaning), the fifteenth (on skepticism), and the sixteenth (on metaphysics).

The excursions can be read in order along with the rest of the book, but they can be read in many different ways, and they can also be skipped as the reader pleases.

I originally intended that this book would contain a ninth chapter on verbal disputes, bringing out a way to use the scrutability framework to help resolve philosophical debates and shedding more light on primitive concepts and the analytic/synthetic distinction. The approach taken there also conveys a more flexible and dynamic version of the framework, without as much philosophical baggage as the earlier chapters. In the end I have left that chapter out: the book is long enough already, and that chapter (which has been published as a separate article) is not quite essential to the overall narrative. Still, I think of that chapter as part of this book in spirit, and it can be found online as part of an extended edition of this book. The same goes for four additional excursions that I have ended up omitting: on inferentialism and analyticity, Twin-Earthability and narrow content, reference magnets and the grounds of intentionality, and conceptual analysis and ordinary language philosophy.

I have been asked a few times what area of philosophy this book falls into. The answer is not obvious, even to me. The book is an unholy stew of epistemology, philosophy of language, metaphysics, and philosophy of mind, with some philosophy of science and metaphilosophy thrown in along the way. But it approaches each of these areas in a distinctive way and with the other areas in mind.

Scrutability theses concern knowledge, so epistemology is at the heart of the project. But the analysis of knowledge, justification, and related notions, which form the core of contemporary epistemology, are only occasionally in focus here. Rather, I am doing a sort of *metaphysical epistemology* (or should that be epistemological metaphysics?): roughly, epistemology in service of a global picture of the world and of our conception thereof.

The metaphysical epistemology in this book breaks down into a number of components. To a first approximation, the early chapters (especially 3 and 4) focus on *global epistemology*: articulating and supporting general theses about what can be known and about the epistemological relations between truths about the world. The middle part of the book (especially chapter 5 and thereabouts) focuses on *epistemological semantics*: understanding notions of meaning and content that are tied to epistemological notions such as rationality and the a priori. The latter part of the book (chapter 7 onward) focuses on *conceptual metaphysics*: roughly, investigating the structure of our conception of reality, with one eye on how well this structure corresponds to reality itself.

The global epistemology in the early chapters serves as the motor that drives the arguments for scrutability for those who are skeptical. Scrutability theses can be seen as global epistemological theses akin to knowability theses and the like. I start by articulating these theses, and then try to argue for them in detail.

Along the way, a lot of epistemology takes place: epistemological issues about warrant, self-doubt, idealization, skepticism, conditionalization, evidence, recognitional capacities, inference, and the a priori take center stage.

The conceptual metaphysics of chapters 7 and 8 serves as the culmination of the book, giving a sense of the full picture that emerges for those who are sympathetic. Here the aim is to boil down our conception of reality to its most basic elements, isolating primitive elements in which our concepts are grounded, and to draw out consequences for mind, language, and reality. The sixteenth excursus draws out the application to issues in metaphysics, fleshing out the projects of conceptual metaphysics and connecting the epistemological notion of scrutability to the related metaphysical notions of supervenience and grounding.

The epistemological semantics of chapter 5 and the excursuses that follow gives a sense of one important application of the framework. Chapter 5 serves to motivate the framework of epistemically possible scenarios and intensions defined over them. The tenth excursus develops the modal framework in more detail. The eleventh excursus develops the semantic framework a little further and argues that the intensions so defined can play many of the key roles of Fregean senses. In chapter 8, I argue that these intensions can serve as a sort of narrow content of thought.

I sketch semantic applications only briefly in this book, but I develop the semantic applications much further in a forthcoming companion volume, *The Multiplicity of Meaning* (and also in various existing articles on which that book is partly based). Where this book starts with Carnap, that book starts with Frege, developing a Fregean approach to language and an internalist approach to thought. There the framework of epistemic two-dimensional semantics, which is itself grounded in the framework of scrutability, plays a central role. The books are written so that either can be read independently of the other, but I think that they work especially well together. They can be read in either order, proceeding either from epistemological foundations to semantic applications or vice versa.

I expect that there will also be a third book at some point, exploring related issues about modality and metaphysics. That book will develop the framework of epistemically possible scenarios, explore its relationship to the space of metaphysically possible worlds, and explore connections to related metaphysical issues. Between them, these three books can be seen as forming a sort of trilogy on the three vertices of the “golden triangle” of reason, meaning, and modality.

The ideas in these books have grown indirectly out of some ideas in my 1996 book *The Conscious Mind*. An early version of the scrutability thesis is explored in chapter 2 of that book, as is a version of the two-dimensional semantic framework that plays a central role in *The Multiplicity of Meaning*. Some of the central themes in the early chapters got their initial airing in the 2001 article “Conceptual Analysis and Reductive Explanation”, co-authored with Frank

Jackson, an article that was at least putatively driven by issues about the mind–body problem.

Despite this connection, it would be a mistake to think of this book as intended mainly to provide a foundation for arguments about the metaphysics of consciousness. If I had been trying to bolster those arguments, I would have written a very different book. In this book, purely metaphysical issues (conceptual metaphysics aside) are most often in the background, while epistemological and semantic issues are in the foreground. In a few places I have articulated theses that might connect the epistemology to the metaphysics (notably the Fundamental Scrutability, Apriority/Necessity, and Conceptual/Metaphysical theses), but I have not tried to argue for them at any length. I have devoted much more energy to arguing for weaker scrutability theses that even thoroughgoing physicalists can accept. The stronger theses and associated metaphysical issues come into focus briefly in chapter 8 and the sixteenth excursus, but they will be more central in the book on modality mentioned above.

It would be somewhat closer to the mark to think of this book as intended to provide a foundation for the ideas about two-dimensional semantics that I have developed in other work. It has gradually become clear to me that the key issue here is scrutability: once an appropriate scrutability thesis is accepted (as I argue in the eleventh excursus), a version of the epistemic two-dimensional framework follows. In fact, this book started its life as a chapter or two in *The Multiplicity of Meaning*, before taking on a life of its own. Still, by now I think that the scrutability thesis has interest for all sorts of purposes, and that while the applications to the theory of meaning and content are important, there are certainly many others as well.

I have tried not to assume too much in the way of theoretical principles from the start. Instead, I have tried to proceed by working through cases and mounting arguments to see what sort of theses emerge at the other side. In this way my approach differs from that of Carnap in the *Aufbau*, who starts with a strong structuralist thesis. I was tempted to write another version of this book, one that first articulates one of the principled scrutability theses in chapter 8 and then uses it to drive a construction from the ground up while also defending it from objections. That principled approach would have been more theoretically elegant and cohesive. But the relatively unprincipled approach of the current book has the advantage of letting the chips fall where they may. This way, by the end of the book we are in a position to judge the prospects for numerous different principled approaches.

Of course I do not proceed with complete philosophical neutrality. There is no such thing, and the discussion here is inevitably filtered through my own philosophical sensibilities. Still, I have tried to acknowledge alternative viewpoints where I can, to find a way for them to come at least part of the way with

me if possible, to argue against them where I can, and to see where they will get off the bus if they must.

My philosophical sensibilities play a role when I consider some of the most famous arguments in recent philosophy: Quine's arguments against analyticity and the a priori, Kripke's arguments against Fregean views, Putnam's and Burge's arguments against internalism. I use the scrutability framework to rebut some of these famous arguments and to limit the consequences of others, thereby defending key elements of the traditional views (internalism, Fregeanism, belief in the a priori, and so on) against which these arguments are directed. I generally take it that the traditional views here have a sort of default status, so that if they are to be rejected it must be on the basis of argument. Because of this way of proceeding, I do not know how much I will do to bring around someone who is entirely unsympathetic with the traditional views, not as a matter of argument but as a matter of starting point. But I am happy enough for now with the conclusion that if these views (or the versions of them that I accept) are wrong, it is for reasons that are interestingly different from the familiar reasons that I argue against.

More generally, it will not surprise me if some of the key conclusions in this book are wrong. Even if ideal reasoners can be certain of the philosophical truth, I am not an ideal reasoner. But I hope that if I am wrong, it is not for old reasons, or not only for old reasons, but also for new and interesting reasons that lead to new and interesting philosophy.

I think of scrutability as supporting a sort of philosophical optimism. Conditional on knowledge of certain fundamental truths and ideal reasoning, everything can be known. In particular, this means that any failures of philosophical knowledge can be ascribed either to the non-ideality of our reasoning or to our ignorance of fundamental truths. Now, it is far from clear to what extent the fundamental truths are knowable, and it is far from clear to what extent we approach the ideal in relevant respects. Still, it is also far from clear that fundamental truths are beyond our grasp, and it is far from clear that reasoning that is needed to determine philosophical truths is beyond our grasp. Philosophy is still young, and the human capacity for reasoning is strong. In a scrutable world, truth may be within reach.

How to Read This Book

It is certainly possible to read this book straight through from start to finish, but the book is long enough that I cannot expect most readers to do that. Some guidance may be helpful, though readers should feel free to ignore it.

Everyone should read the introduction and chapter 1, which introduce the project, and at least browse chapter 8, which discusses conclusions and some applications. For a minimal path in between, it is possible to read just the first section or two of each of the intervening chapters, especially chapters 2, 3, and 7. The remainder of chapters 2–4 formulate and argue for scrutability theses in considerable detail, and it is quite possible to skip this detail on a first reading, although readers who are skeptical about scrutability theses might want to focus here. Chapters 6 and 7 (and to some extent 8) focus on the question of just what needs to be in a scrutability base. Chapter 7 is the more important of the two (and provides useful background to chapter 8), but readers should feel free to skip to the cases that are most interesting to them. Chapter 5 is not essential to the central narrative, but it goes into important foundational issues about the a priori while also developing the application of the framework to the analysis of meaning. Everyone should read the summation after chapter 8, which sums up some main conclusions.

Readers should feel free to dip into the excursuses as they like. Many of them can be read on their own, at least given chapter 1 as background, although there is usually some connection to the preceding chapter. A few of the excursuses (especially 3–7 and 14) go into details that are in principle essential to the theses and arguments of the book, but that in practice can be skipped by those who are not too concerned with the relevant issues. Some others (1–2, 8–9, 13) develop connections to related literature. Further excursuses (especially 10–12 and 15–16) develop applications of the framework to issues about modality, meaning, the unity of science, skepticism, and metaphysics. A number of further applications are outlined in chapter 8.

The six John Locke lectures correspond roughly to portions of chapters 1 (along with the introduction), 3, 4, 5, 6, and 8 (along with the summation) respectively, with some elements of chapters 2 and 7 thrown in along the way, and with almost none of the material from the excursuses.

The extended edition, containing an additional chapter (on verbal disputes) and four additional excursuses (on inferentialism, reference magnets and the grounds of intentionality, conceptual analysis and ordinary language philosophy, and Twin-Earthability) can be found online by searching for “Constructing

the World: Extended Edition”. I have developed a number of relevant points further in various articles, usually cited in the text with just an article title. All of these articles are also available online.

Readers without a background in philosophy will probably find this book hard going, but some relatively accessible material includes the introduction and chapter 1, the material on the Cosmoscope in chapter 3, the first half of chapter 5, and intermittent sections of chapters 6 and 7. A glossary at the end of the book provides a guide both to expressions and theses that I have introduced and to some commonly used philosophical expressions that I deploy.

Issues in many areas of philosophy are discussed in this book. Here I have indicated where to find material especially relevant to certain areas. In what follows “1” stands for chapter 1, “1.1” stands for section 1 of chapter 1, “1.1–2” stands for sections 1 and 2 of chapter 1, “E1” stands for the first excursus, and so on. Some of these abbreviation formats (“1.1”, “E1”) are also used for cross-referencing throughout the book.

Informal epistemology: 2.3, 3.3–4, 3.7, 4, 8.5, E1, E4, E7, E8, E15.

Formal epistemology: 2.4, 4.2, 5.6, 5.7, E5.

Metaphysics and modality: 1.5, 6.4, 6.6, 6.16, 7.3, 7.9, 8.6, E6, E10, E16.

Philosophy of language: 1.1, 1.3, 1.4, 2.2, 5, 6.9–13, E2, E3, E9, E11, E14.

Philosophy of mind: 1.1, 3.7, 6.7, 6.14, 7.4–5, 7.7, 8.3–7, E9, E17.

Philosophy of science: 6.2, 6.15, 7.3, 7.5, 7.8, 8.7, E9, E12.

Carnap and logical empiricism: 1, 3.1, 5, 8, E1, E10, E12, E15.



1

Scrutability and the *Aufbau*

I Primitive concepts

What are the basic elements of thought? It is common to hold that thoughts, such as *Galabs are pink*, are composed of concepts, such as *galah* and *pink*. It is also common to hold that many concepts are composed from simpler concepts. For example, Aristotle held that ‘man’ can be defined as ‘rational animal’. This suggests that the concept *man* is a complex concept built out of the simpler concepts *rational* and *animal*.

In his manuscript “De Alphabeto Cogitationum Humanarum”, Leibniz suggests that there is a level of concepts so simple that they make up an alphabet from which all thoughts can be composed:

The alphabet of human thoughts is a catalog of primitive concepts, that is, of those things that we cannot reduce to any clearer definitions.¹

In *An Essay Concerning Human Understanding*, John Locke develops such a picture. He introduces complex ideas (or concepts) as follows:

As simple ideas are observed to exist in several combinations united together, so the mind has a power to consider several of them united together as one idea; and that not only as they are united in external objects, but as itself has joined them together. Ideas thus made up of several simple ones put together, I call complex;—such as are beauty, gratitude, a man, an army, the universe. (Locke 1690, book 2, chapter 12)

Locke held that all of our perception and thought derives from simple ideas. At one point in the *Essay* (book 2, chapter 21), he suggests that the most basic ideas come down to eight. Three are ideas of matter that come to us through our senses: *extension*, *solidity*, and *mobility* (the power of being moved). Two are ideas

¹ Translated from “De Alphabeto Cogitationum Humanarum” (A 6.4.270), written around 1679–81. Thanks to Brandon Look for the translation.

of mind that come to us through reflection: *perceptivity* (the power of perception or thinking) and *motivity* (the power of moving). The last three are neutral ideas that come to us both ways: *existence*, *duration*, and *number*.

The same theme can be found in some parts of contemporary cognitive science. The linguist Anna Wierzbicka, for example, has argued that every expression in every human language can be analyzed in terms of a limited number of “semantic primes” that occur in every language. In her 1972 book *Semantic Primitives*, Wierzbicka proposed 14 semantic primes, but by her 2009 book *Experience, Evidence, and Sense* these had expanded to the following list of 63 primes.

Substantives: *I, you, someone, something/thing, people, body.*

Relation substantives: *kind, part*

Determiners: *this, the same, other/else.*

Quantifiers: *one, two, some, all, much/many.*

Evaluators: *good, bad.*

Descriptors: *big, small.*

Mental predicates: *think, know, want, feel, see, hear.*

Speech: *say, words, true.*

Actions and events: *do, happen, move, touch.*

Existence, possession, location: *to be (somewhere), there is/there are, to have, to be (someone/something).*

Life and death: *live, die.*

Time: *when/time, now, before, after, a long time, a short time, for some time, in a moment.*

Space: *where/place, here, above, below, far, near, side, inside.*

Logic: *not, maybe, can, because, if.*

Augmentors: *very, more.*

Similarity: *like.*

Wierzbicka’s methods have been used to analyze an extraordinary range of expressions in many different languages. To give the flavor of the project, a sample analysis runs as follows.

X *lied* to Y =

X said something to person Y;

X knew it was not true;

X said it because X wanted Y to think it was true;

people think it is bad if someone does this.

In twentieth-century philosophy, this sort of framework was developed most systematically by Bertrand Russell and Rudolf Carnap.² Russell suggested that

² Russell engaged in numerous different projects of analysis and construction. Some central works concerning analysis into primitives involving acquaintance include “Knowledge by Acquaintance and Knowledge by Description” (1911), *The Problems of Philosophy* (1912), and “Theory of Knowledge” (1913). He pursued related projects of constructing the world from primitives

all concepts are composed from concepts of objects and properties with which we are directly acquainted. For Russell, these concepts included concepts of sense-data and certain universals, and at certain points in his writings, a concept of oneself. All other concepts were to be analyzed as constructions out of these concepts. For example, concepts of other people and of objects in the external world were to be analyzed as descriptions built up from these basic elements.

In the *Der Logische Aufbau der Welt*, Carnap pushed the project of analysis to its limit. Carnap argued that all concepts can be constructed from a single primitive concept, along with logical concepts. Carnap's primitive concept was a concept of the relation of phenomenal similarity: similarity in some respect between total experiences (roughly, momentary slices of a stream of consciousness) had by a subject at different times.³ For example, if a subject has two experiences both involving a certain shade of red, the experiences will stand in this relation of similarity. Using this simple concept, Carnap gave explicit constructions of many other concepts applying to experiences. For example, concepts of specific sensory qualities, such as that of a certain shade of red, are defined in terms of chains or circles of similarity between experiences.

In Carnap's framework, these concepts are used to build up all of our concepts of the external world. Spatial and temporal concepts are defined in terms of sensory qualities. Properties of external bodies are defined in terms of spatial and temporal properties. Behavior is defined in terms of the motion of bodies. Mental states of other people are defined in terms of behavior. Cultural notions are defined in terms of these mental states and behavior. And so on.⁴

Carnap's project, like most of the other projects above, is committed to what we can call a *Definability* thesis. Like the other theses I discuss in this chapter, this thesis is cast in terms of expressions (linguistic items such as words) rather than in terms of concepts (mental or abstract items) for concreteness.

Definability. There is a compact class of primitive expressions such that all expressions are definable in terms of that class.

I will say more about compactness later, but for now we can think of this as requiring a small class of expressions. For most of the *Aufbau*, the class of primitive expressions included an expression for phenomenal similarity and logical

in numerous later works, such as "The Philosophy of Logical Atomism" (1918). Also worth mentioning is Ludwig Wittgenstein's conception of the world as the totality of atomic facts in his *Tractatus Logico-Philosophicus* (1921), although Wittgenstein says less than Carnap and Russell about the character of his primitives and about the construction of ordinary concepts.

³ In this book 'phenomenal' always means "experiential": roughly, pertaining to conscious experiences.

⁴ It must be acknowledged that the details are sometimes sketchy. See the start of chapter 6 for an illustration of Carnap's treatment of culture.

expressions ('not', 'and', 'exists', and the like). Late in the *Aufbau*, Carnap went on to argue that phenomenal similarity is itself dispensable: it can itself be defined in logical terms. If so, then primitive expressions can be restricted to logical expressions, and all other expressions can be defined in terms of these. Of course the general program of definability is not committed to as strong a claim as this.

We can say that an expression *E* is definable in terms of a class of expressions *C* if there is an adequate definition statement with *E* on the left-hand side and only expressions in *C* on the right hand side. We then need to say what a definition statement is, and what it is for such a statement to be adequate.

A definition statement connects a left-hand side involving a defined expression *E* to a right-hand side, with a logical form that depends on the grammatical category of *E*. Various different logical forms might be required, but the differences will not matter for our purposes. As an example, definition statements for singular terms, general terms, and predicates might be required to specify the extension of *E* (roughly, the entity or entities in the world that *E* applies to) in a form akin to the following: 'For all *x*, *x* is Hesperus if and only if *x* is the brightest object visible in the evening sky'; 'For all *x*, *x* is a bachelor if and only if *x* is an unmarried man'. If such definition statements are adequate, then 'Hesperus' is definable in terms of 'brightest', 'evening', and so on, and 'bachelor' is definable in terms of 'unmarried', 'man', and so on.

What is it for a definition statement to be adequate? Here, there are various possible criteria. Certainly one should require at least *extensional* adequacy: that is, definitions of the sort above must be true, so that the extensions of the relevant expressions on the left and right sides are the same. But typically more is required. Suppose that as it happens, all bachelors in our world are untidy men and vice versa. Then 'For all *x*, *x* is a bachelor if and only if *x* is an untidy man' is true, and the definition statement is extensionally adequate. Still, this statement does not seem to give an adequate definition of 'bachelor'.

To handle these cases, it is common to require some form of stronger-than-extensional, or *intensional*, adequacy for a definition. For example, it is often required that a definition statement be analytic (true in virtue of meaning), a priori (knowable without justification from experience), and/or necessary (true in all possible worlds). A definition of 'bachelor' in terms of 'untidy man' does not meet these conditions: 'all bachelors are untidy men' is not true in virtue of meaning, one cannot know a priori that all bachelors are untidy men, and it is not true in all possible worlds that all bachelors are untidy men. But it is at least arguable that a definition of 'bachelor' in terms of 'unmarried man' meets these conditions.

A surprising and often-overlooked feature of the *Aufbau* is that Carnap there requires only that definitions be extensionally adequate. Carnap intended the *Aufbau* to shed light on knowledge and on meaning, but it is questionable whether definitions that are merely extensionally adequate can fulfill these epistemological and semantic goals. For example, while a definition of 'bachelor' as

‘unmarried man’ may shed some light on the meaning of ‘bachelor’ and on how we come to know truths about bachelors, the same does not seem true of a definition as ‘untidy man’, even if that definition is extensionally adequate. In the preface to the second edition of the *Aufbau*, Carnap says that this is the greatest mistake in the project, and says that definitions should be held to a stronger, intensional, criterion of adequacy. Certainly much of the *Aufbau* can be read with a stronger criterion of adequacy in mind.⁵

The stronger criteria of analyticity, apriority, and necessity ensure that an expression and its definition are connected semantically (that is, in the realm of meaning), epistemologically (in the realm of knowledge), and modally (in the realm of necessity and possibility). Further potential criteria include psychological criteria, to the effect that a definition somehow reflects the psychological processes involved in understanding and using an expression; formal criteria, to the effect that definitions have a certain limited complexity; conceptual criteria, to the effect that the expressions used in the definition express concepts that are more basic (in some relevant sense) than the concept expressed by the original expression; and so on.

Definitions allow us to connect sentences in different vocabularies. Given a definition of bachelors as unmarried men, truths such as ‘John is a bachelor’ will be logically entailed by truths such as ‘John is an unmarried man’ along with the definition. More generally, given certain assumptions about the language,⁶ any statement containing ‘bachelor’ will be logically entailed by a corresponding sentence containing ‘unmarried man’ in place of ‘bachelor’ (with the rest of the sentence as before), along with the definition. Given these assumptions, the Definability Thesis leads to the following thesis:

Definitional Scrutability: There is a compact class of truths from which all truths are definitionally scrutable.

⁵ Carnap sometimes explicitly invokes stronger criteria in the *Aufbau*. For example (as Chris Pincock pointed out to me), in section 49 he suggests a method according to which constructional definitions for scientific objects are determined by their epistemological “indicators”.

⁶ We can assume that every natural-language sentence has a *regimentation* into an equivalent sentence with a clarified logical form. One can then apply definitional and logical machinery to regimented sentences in the first instance, and derivatively to unregimented sentences. If definitions are required only to be extensionally adequate, it suffices to assume that the language and the logic are extensional: that is, the logic allows one to substitute coextensive expressions (given a statement of coextensiveness), and this substitution will not change truth-values in the language. If definitions are required to be intensionally adequate, it suffices to assume that the language and the logic are intensional to the same degree (definitions will then need to contain a statement of cointensiveness, such as ‘Necessarily, bachelors are unmarried men’). The language may also be hyperintensional, so that cointensive expressions are not intersubstitutable in certain contexts, as long as these contexts can themselves be defined in an extensionally/intensionally adequate way (‘For all S, S believes that such-and-such if...’).

Here, a truth is a true sentence.⁷ A compact class of truths, to a first approximation, is a class of truths involving only a small class of expressions. A sentence S is definitionally scrutable from (or definitionally entailed by) a class of sentences C if S can be logically derived from some members of C and some adequate definition sentences. For example, given the relevant assumptions, sentences involving ‘bachelor’ are definitionally scrutable from sentences involving ‘unmarried man’. If we repeat this process for every definable expression, we can eventually translate every sentence of the language into a sentence in the primitive vocabulary, and the original statement will be entailed by the transformed sentence conjoined with a number of definitions.

On the *Aufbau* view, all truths are definitionally scrutable from a class of truths about the phenomenal similarity relation. In fact, Carnap holds that there is a single *world-sentence* D that definitionally entails all truths. The world-sentence says that there exist entities that are related in such-and-such fashion by the phenomenal similarity relation R . If there are just two dissimilar total experiences in the world, then the world-sentence will be a sentence saying that there are two entities that stand in R to themselves but not to each other: $\exists x, y : (Rxx \& Ryy \& \sim Rxy \& \sim Ryx \& \forall w (w = x \vee w = y) \& \sim (x = y))$. If there are more total experiences than this, then there will be a longer world sentence, specifying the similarity relations that do and do not hold among the total experiences.

According to Carnap’s stronger view late in the *Aufbau*, the previous world-sentence D is definitionally entailed by an even more austere world sentence D' , using purely logical vocabulary. To get from D to D' , Carnap defines away the single nonlogical vocabulary item R as that relation that makes the previous world-sentence D true.⁸ If this is correct, then the highly austere truth D' definitionally entails all truths.

If we require that adequate definitions are a priori (knowable independently of experience), as is common, then Definitional Scrutability entails the following thesis:

A Priori Scrutability: There is a compact class of truths from which all truths are a priori scrutable.⁹

⁷ The choice of sentences rather than propositions here is discussed in 2.2. A subtlety here (discussed at length in the third excursus) is that not all sentences are true or false independent of context. For example, there may be no context-independent fact of the matter about whether a sentence such as ‘I am hungry’ or ‘John is tall’ is true. Where context-dependent sentences are concerned, we can talk instead of the scrutability of sentences in contexts.

⁸ For the world-sentence just specified, R will be defined as that relation R' such that: $\exists x, y : (R'xx \& R'yy \& \sim R'xy \& \sim R'yx \& \forall w (w = x \vee w = y) \& \sim (x = y))$. Then the new world-sentence will be the resulting of replacing R in the world-sentence above by this definition. Or more straightforwardly, the world-sentence can simply say $\exists R', x, y : (R'xx \& R'yy \& \sim R'xy \& \sim R'yx \& \forall w (w = x \vee w = y) \& \sim (x = y))$.

We can define a priori scrutability in parallel to definitional entailment: a sentence S is a priori scrutable from (or a priori entailed by) a class of sentences C if S can be logically derived from some members of C along with some a priori truths. Given weak assumptions,⁹ the right-hand side is equivalent to the claim that there is a conjunction D of sentences in C such that the material conditional 'If D , then S ' is a priori.

One can characterize Analytic and Necessary Scrutability theses in a parallel way. If we require that adequate definitions are analytic or necessary respectively, then these theses will follow from Definitional Scrutability.

It is theses such as A Priori and Analytic Scrutability that give the definitional program its epistemological bite. To a first approximation, these theses suggest that knowledge of the base truths about the world might serve as a basis for knowledge of all truths about the world.

To make this vivid: suppose that Laplace's demon is given all the base truths about our world. Given Definitional Scrutability, then as long as the demon knows all the definitions and can engage in arbitrary logical reasoning, then the demon will be able to deduce all truths about the world. Given A Priori Scrutability, then as long as the demon can engage in arbitrary a priori reasoning, then it will be able to deduce all truths about the world. For example, if Carnap is right, then the demon should be able to derive all truths about the world from a world sentence such as D or D' .

2 Objections to the *Aufbau*

The *Aufbau* is widely held to be a failure. It is also widely held that no project like it can succeed. These doubts have a number of sources. Perhaps the best-known problems for the *Aufbau* are arguments that Carnap's primitive vocabulary cannot do the work it needs to do. Two of these are specific criticisms of Carnap's constructions from phenomenal vocabulary, while another two are general criticisms of constructions from phenomenal vocabulary or from logical vocabulary.

First: In *The Structure of Appearance* (1951), Nelson Goodman argued that Carnap's definition of sensory qualities in terms of the primitive of recollected

⁹ A more elaborate definition of a priori scrutability is given in 2.5, and a more elaborate discussion of what it is for a sentence to be a priori is in 4.1.

¹⁰ In one direction, it suffices to assume that all conjunctions are logically derivable from their conjuncts (this is trivial in the finite case, but slightly less trivial if infinite conjunctions are allowed, as may be necessary for some purposes). In the other direction, it suffices to assume that when B is logically derivable from a set A of premises, a conditional 'If D then B ' is a priori, where D is a conjunction of the premises in A , and that a priori conjuncts can be detached from the antecedents of a priori conditionals without loss of apriority.

phenomenal similarity is unsuccessful, as there can be circles of similarity among total experiences that do not correspond to a single sensory quality. One problem raised by Goodman is that of “imperfect community”: a similarity circle can satisfy Carnap’s definition of a sensory quality even when some members of the circle share one quality (phenomenal redness, say) and others share another quality (phenomenal blueness). Another problem is that of “companionship”: if two distinct qualities always occur together in total experiences, Carnap’s definition will not distinguish them.

Second: In “Two Dogmas of Empiricism” (1951), W. V. Quine argued that Carnap’s definition of spacetime points in terms of the phenomenal field is unsuccessful, as it requires nonphenomenal notions that violate his own criteria of adequacy. Carnap defined “Quality q is at x, y, z, t ” by specifying certain principles for assigning qualities to spacetime points that must be obeyed as well as possible, but this does not yield a definition that can be cast entirely in terms of phenomenal notions and logic.

Third: In “The Problem of Empiricism” (1948), Roderick Chisholm gives a general argument against phenomenalism: the view that statements about the external world can be definitionally analyzed in purely phenomenal terms. On a phenomenalist view, ‘There is a doorknob in front of me’ (P) must be analyzed as a complex conditional along the lines of ‘If I had certain experiences, certain other experiences would follow’ (R): for example, ‘If I experience a certain sort of attempt to grasp, I would experience a certain sort of contact’. Chisholm argues that no such R is entailed by P , as one can always find a further sentence S (e.g. specifying that one is paralyzed and subject to certain sorts of delusions of grasping that are never accompanied by experiences of contact) that is consistent with P such that $S \& P$ entails $\neg R$. If so, no phenomenalist analysis of P can succeed.

Fourth: In “Mr. Russell’s Causal Theory of Perception” (1928), the mathematician Max Newman pointed out a general problem for the more ambitious project of reducing the primitive vocabulary to logical structure alone. The problem was pointed out simultaneously by Carnap himself late in the *Aufbau*.¹¹ Given a purely logical vocabulary, the ultimate world-sentence (like D' above) will specify simply that there exist certain objects, properties, and relations that stand in certain patterns of instantiation and co-instantiation. Newman and Carnap observe that as long as we are liberal enough about what we count as a property or a relation, this world-sentence will be satisfied almost vacuously.¹² Carnap responds by suggesting that the properties and relations in question must be

¹¹ Carnap marks these sections of the *Aufbau* (153–55) “can be omitted”, quite remarkably given the centrality of these sections to the “logical structure” project. For further discussion of Newman’s problem and the *Aufbau*, see Demopolous and Friedman 1985.

restricted to “natural” (or “founded”, or “experientable”) properties and relations. This requires an expansion of the primitive vocabulary, which Carnap justifies by suggesting that ‘natural’ is a logical term. Few have found this latter suggestion convincing, however.

Still, it is clear that criticisms of this sort threaten only *Aufbau*-style projects that involve phenomenal and/or logical bases. To avoid the problems, one need only expand the primitive basis. One can avoid Newman’s problem by allowing almost any nonlogical vocabulary. One can avoid Goodman’s problem by allowing expressions for specific sensory qualities. One can avoid Quine’s and Chisholm’s problems by allowing spatiotemporal expressions into the basic vocabulary directly, or perhaps by allowing expressions for causal relations.¹³

One might wonder whether expanding the base like this is in the spirit of the *Aufbau*. For many years, the popular conception of logical empiricism has focused on a commitment to phenomenalism and verificationism (views on which a phenomenal base is central), and the *Aufbau* has been regarded as a paradigm of that tradition.¹⁴ In reality, these views do not play a central role in the *Aufbau*. A much more important role is played by Carnap’s commitment to structuralism and objectivity in developing a language for science. Carnap himself says that the choice of a phenomenal basis in the *Aufbau* is somewhat arbitrary, and that he could equally have started with a physical basis. A base with expressions for specific sensory qualities or specific physical properties (such as

¹² In particular, as long as there is a property corresponding to any set of objects, and a relation corresponding to any set of ordered pairs, then the world-sentence S will be satisfied by any set of the right size. To see this, suppose that one set A of size n satisfies S , and let A' be any other set with the same size. Take a group of properties and relations that relate the members of A in the pattern specified by S . Map those properties and relations to a corresponding set of properties and relations on A' by a one-to-one mapping. (Any one-to-one mapping will do; the liberalness claim will ensure that every property maps to a property, and so on.) Then the resulting properties and relations will relate the members of A' in the same pattern. So S will be satisfied by A . It follows that S cannot entail any truths that specify features of the world beyond its cardinality.

¹³ Even while retaining a phenomenal base, Carnap has some options in avoiding the first three problems. Carnap’s construction is defended against Goodman and Quine by Thomas Mormann (2003, 2004), while a different construction from an expanded phenomenal base is explored by Hannes Leitgeb (2011).

¹⁴ The distortions in the popular conception of the *Aufbau* and logical empiricism are explained partly by simplified versions promulgated by A. J. Ayer and W. V. Quine, and partly by a post-*Aufbau* period in the Vienna Circle in which phenomenal reductions involving protocol sentences played a more crucial role. Within a few years of that period (for example, in his 1932 work “The physical language as the universal language of science”) Carnap had moved on again to a view on which physical language rather than phenomenal language plays the crucial role in reduction. In recent years, the flourishing scholarly literature on the *Aufbau* and logical empiricism, including Alberto Coffa’s *The Semantic Tradition from Kant to Carnap* (1985), Michael Friedman’s *Reconsidering Logical Positivism* (1999), Alan Richardson’s *Carnap’s Construction of the World* (1998), and Thomas Uebel’s *Overcoming Logical Positivism from Within* (1992), among other works, has painted a picture that is much more nuanced than the popular caricature.

spatiotemporal properties) might not fully vindicate Carnap's structuralism, but as I discuss in chapter 8, there are other bases that come even closer to fulfilling Carnap's goals. In any case, expanded bases have the potential to fulfill many of the more general aims of a project of definability, while avoiding the criticisms above.

Other doubts about the project of the *Aufbau* are driven not by Carnap's basic vocabulary but by his construction method: that is, by his method of deriving nonbasic truths from basic truths using definitions. A number of doubts about definitions have been influential.

First: In "Verifiability" (1945), Friedrich Waismann argued that purported definitions of ordinary expressions are subject to the problem of *open texture*: these definitions are always subject to correction, as we cannot foresee all possibilities to which they might apply. Every definition "stretches into an open horizon", and no definition of an empirical term will cover all possibilities. Waismann's argument was especially directed at definitions in the style of logical empiricism that appeal to methods of verification, but his underlying point applies quite generally.

Second: In the *Philosophical Investigations* (1953), Ludwig Wittgenstein suggested that when we apply a term such as 'game' to some things, there is no single condition that they all satisfy. 'Game' is a family resemblance term, with different sorts of games resembling each other in various respects and with no common core. There is merely a "complicated network of similarities, overlapping and criss-crossing". Many have taken this idea to suggest that there are no definitions giving necessary and sufficient conditions associated with ordinary expressions of this sort.

Third: In "Two Dogmas of Empiricism" (1951), Quine gave a critique of the notion of definition and more generally of the analytic/synthetic distinction. He argued that standard understandings of these notions are circular and that the notions are based on a misconceived picture of language and its relation to the world. This critique has led many to doubt that a substantial distinction between the analytic and the synthetic, or between the a priori and the a posteriori, or between the definitional and the nondefinitional, can be drawn. If these doubts are correct, then any *Aufbau*-like project that involves these notions must fail.

Fourth: In *Naming and Necessity* (1980), Saul Kripke argued against descriptivism: the thesis that names are equivalent to descriptions. Kripke's modal argument makes a case that for an ordinary name (e.g. 'Aristotle') and an associated description (e.g. 'the teacher of Alexander'), the name and the description are not necessarily equivalent. Kripke's epistemic argument makes a case that for an ordinary name (e.g. 'Gödel') and an associated description (e.g. 'the man who proved the incompleteness of arithmetic'), the name and the description are not a priori equivalent. If these arguments succeed, then it appears that no *Aufbau*-

like definitional project that applies to names and that invokes necessity or apriority as a condition of adequacy can succeed.

These criticisms mainly threaten an *Aufbau*-style project whose construction relation requires definitions of nonbasic expressions. Just as we can get around the first class of problems by expanding the base, we can get around the second class of problems by weakening the construction relation.

Before doing that, however, it is useful to look more closely at the source of the problems. At least three of the critiques (Waismann's, Wittgenstein's, and Kripke's) turn on a common problem: the problem of *counterexamples*. (Quine's critique turns on somewhat different issues, and I return to it in Chapter 5.) For many terms in English, it seems that every definition that has ever been offered is subject to counterexamples: actual or possible cases to which the original term applies but the purported definition does not, or vice versa, thereby showing that the definition is inadequate.

The most famous case is the case of 'knowledge', traditionally defined as 'justified true belief'. In his 1963 paper "Is Knowledge Justified True Belief?", Edmund Gettier pointed out counterexamples to this purported definition. Suppose that Smith has a justified belief that Jones owns a Ford, and deduces that Jones owns a Ford or Brown is in Barcelona. And let us say that Jones has recently sold his Ford, and that Brown is in fact in Barcelona, though Smith has no information about either of these things. Then Smith has a justified true belief that Jones owns a Ford or Brown is in Barcelona, but this justified true belief is not knowledge. So knowledge cannot be defined as justified true belief.

In Gettier's wake, others attempted to modify the definition of knowledge to avoid these counterexamples, for example suggesting that knowledge can be defined as justified true belief that is not essentially grounded in a falsehood. But other counterexamples ensued: if I see the one real barn in an area of fake barns, and form the belief that I am seeing a barn, then this is a justified true belief not essentially grounded in a falsehood, but it is not knowledge. A parade of further attempted definitions and further counterexamples followed (Shope's *The Analysis of Knowing* gives an exhaustive summary). Eventually definitions with fourteen separate clauses were proffered, with no end to the counterexamples in sight.

What goes for 'knowledge' seems to go for most expressions in the English language. Given any purported definition of 'chair', or 'run', or 'happy', it is easy to find counterexamples. For some scientific terms such as 'gold' or 'electron', there may be true definition statements ('Gold is the element with atomic number 79'), but these do not appear to be a priori. For Wierzbicka's definition of 'lie', above, counterexamples are not hard to find: I can tell a lie even if I do not care whether the hearer believes me.¹⁵ And even in the case of 'bachelor', there are unmarried men who do not seem to be bachelors, such as those in

long-term domestic partnerships. The only clearly definable expressions appear to be derived expressions (such as ‘unhappy’ and ‘caught’), which can arguably be defined in terms of the expressions (‘happy’ and ‘catch’) that they are derived from, along with some technical expressions that have been introduced through definitions, and a handful of others.

The philosophical flight from definitions has been paralleled by a similar flight in cognitive science. Contemporary psychologists almost universally reject the so-called classical view of concepts, according to which most concepts are associated with sets of necessary and sufficient conditions. A major influence here is work by Eleanor Rosch (1975) and others on concepts such as that of a bird, suggesting that subjects classify various creatures as birds in a graded way according to their similarity to various prototypes rather than by necessary and sufficient conditions.¹⁶ By and large, the classical view has been supplanted by views on which concepts involve prototypes, exemplars, and theories, among other views. On few of these views is it required that concepts are associated with definitions.

It remains possible that for these expressions, there exists an adequate definition that has not yet been found. In philosophy, the search for definitions typically runs out of steam once purported definitions reach a certain length. In psychology, it is not out of the question that prototype theories and the like might be used to deliver something like a definition, perhaps cast in terms of weighted similarities to certain prototypes or exemplars. Likewise, theory-based accounts of concepts might yield definitions of various concepts in terms of clusters of associated theoretical roles. Still, it is far from obvious that such definitions will exist, and even if they do exist, they will be so unwieldy that they will be quite unlike definitions as traditionally conceived. As a result, the definitional program has been put to one side in most areas of philosophy and psychology in recent years.

3 From Definitional Scrutability to A Priori Scrutability¹⁷

Even if Definitional Scrutability is false, there remains a strong case for other scrutability theses. For example, even if expressions such as ‘knowledge’ and ‘chair’ are not definable in terms of more primitive expressions, it remains plausible that there is some strong epistemological relation between truths

¹⁵ A philosopher will find possible counterexamples to many or most of Wierzbicka’s definitions. Wierzbicka’s intended criteria of adequacy for definitions almost certainly differ from philosophers’ criteria, so it is not obvious to what extent the existence of counterexamples is a problem for Wierzbicka’s project.

¹⁶ A distinct anti-definition influence in psychology derives from psycholinguistic arguments for the conclusion that lexical concepts are primitive by Jerry Fodor et al. (1980).

¹⁷ This section overlaps in part with Chalmers and Jackson 2001.

involving these expressions and truths involving more primitive expressions. In particular, it is striking that in many cases, specifying a situation in terms of expressions that do not include ‘knowledge’ or its cognates (synonyms or near-synonyms) enables us to determine whether or not the case involves knowledge. Likewise, correctly describing an object in terms of expressions that do not include ‘chair’ or its cognates may enable us to determine whether or not it is a chair. And so on.

For example, in the Gettier situation we are told something like:

‘Smith believes with justification that Jones owns a Ford. Smith also believes that Jones owns a Ford or Brown is in Barcelona, where this belief is based solely on a valid inference from his belief that Jones owns a Ford. Jones does not own a Ford, but as it happens, Brown is in Barcelona.’

Let the conjunction of these sentences be G . G does not contain the term ‘know’ or any cognates. But when presented with G , we are then in a position to determine that the following sentence K is false:

‘Smith knows that Jones owns a Ford or Brown is in Barcelona.’

Something like this happens throughout philosophy, psychology, and other areas. We are given a description D of a scenario without using a key term E , and we are asked to determine whether and how the expression E applies to it. This is the key method for experimental work on concepts in psychology: an experimenter presents a description (or perhaps a picture) of a case, subjects are asked to classify it under a concept, and they usually can do so. The same goes for conceptual analysis in philosophy: one considers a specific case, considers the question of whether it is a case of an F , and one comes to a judgment. Often we have no trouble doing so.

In fact, this method of cases is precisely how counterexamples to definitions are often generated. When someone suggests that E can be defined as F (‘bachelor’ is defined as ‘unmarried man’, say), someone else suggests a scenario D (involving long-term gay couples, say) to which F applies but E does not, or vice versa. The Gettier case fits this pattern perfectly. Despite the absence of definitions, there is some form of scrutability present in these cases: once we know G , we are in a position to know $\sim K$, and so on.

In many cases, it is plausible that the scrutability is a priori. For example, in the Gettier case, it is plausible that one can know the material conditional ‘If G , then $\sim K$ ’ a priori. Someone who knows that G is true and who has mastered the concepts involved in K (in particular the concept of knowledge) is thereby in a position to know that K is false, even if they lack any further relevant empirical information. That is, mastery of the concept of knowledge (along with a grasp of

the other concepts involved) and rational reflection suffices to eliminate the possibility that both G and K are true.

On the face of things, Gettier's argument was an a priori argument, in which empirical information played no essential role, and its conclusion is a paradigmatic example of a non-obvious a priori truth. The argument proceeds by presenting the hypothesis that G holds, and appealing to the reader's possession of the concept of knowledge to make the case that if G holds, $\sim K$ holds (and J holds, where J is a corresponding positive claim about Smith's justified true belief). Empirical information plays no essential role in justifying belief in this conditional, so the conditional is a priori. The a priori conditional itself plays an essential role in deriving the a priori conclusion.

This brings out a key point: a priori scrutability does not require definability. One might think that for a sentence B to be a priori entailed by a sentence A , the terms in B must be definable using the terms of A . However, this thesis is false. The a priori entailment from 'There exists a red ball' to 'There exists a colored ball' is one counterexample: 'colored' cannot be defined in terms of 'red' and the other terms involved. But the case above is another counterexample. At least once general skepticism about the a priori is set aside, 'If G then $\sim K$ ' is a central example of an a priori truth. But at the same time, we have seen that there is little reason to think that there is an adequate definition of 'knowledge', whether in the terms involved in G or any other terms.

As before, it could be that there is an adequate definition that has not yet been produced, or that has been produced but overlooked. Someone might even hold that all these a priori conditionals are underwritten by our tacit grasp of such a definition. But even if so, it seems clear that the a priori entailment from G to $\sim K$ is not dialectically hostage to an explicit analysis of knowledge that would support the entailment. That is, we can have reason to accept that there is an a priori entailment here even without having reason to accept that there is an explicit analysis that supports the entailment.

If anything, the moral of the Gettier discussion is the reverse: at least dialectically, the success of a definition itself depends on a priori judgments concerning specific cases, or equivalently, on a priori judgments about certain conditionals. The Gettier literature shows repeatedly that purported definitions are hostage to specific counterexamples, where these counterexamples involve a priori judgments about hypothetical cases. So a priori conditionals seem to be prior to definitions at least in matters of explicit justification. Our judgments about a priori conditionals do not need judgments about definitions to justify them, and are not undermined by the absence of definitions.

It might be suggested that our conditional judgments here require at least explicit *sufficient* conditions for knowledge or its absence: for example, the condition that a belief based solely on inference from a false belief is not knowl-

edge. It is trivial that there is a sufficient condition in the vicinity of such an entailment (the antecedent provides one such), so the claim will be interesting only if the complete set of sufficient conditions for knowledge is not huge and open-ended. But the Gettier literature suggests precisely that the set of sufficient conditions for knowledge is open-ended in this way; if it were not, we would have a satisfactory definition. And as before, the a priori entailments are not dialectically hostage to the proposed sufficient conditions. Rather, at least in common practice, proposed sufficient conditions are hostage to a priori intuitions about specific cases.

It may even be that there are no short nontrivial sufficient conditions for knowledge. That is, it may be that any reasonably short condition not involving ‘know’ or cognates is compatible with the absence of knowledge.¹⁸ Not every expression is like this. For example, there are plausibly short sufficient conditions for *not* knowing that *p*: the condition of not believing that *p*, or of believing that *p* based solely on inference from a false belief. But it may be that for many expressions, there are at least hypothetical cases for which there is no reasonably short nontrivial sufficient condition (perhaps even no finite sufficient condition) obtaining in that case. In such a case, a nontrivial sufficient condition must be a long one: in the limit, a fully detailed specification of such a scenario, perhaps in the language of a scrutability base. All this is quite consistent with A Priori Scrutability, but it does bring out the need for idealization in understanding the thesis.

An opponent of A Priori Scrutability may hold that there are not even long nontrivial sufficient conditions for knowledge and the like, or that any sufficient conditions here do not yield a priori scrutability. These remain separate substantive issues, distinct from the standard objections to Definability and addressed in the arguments for A Priori Scrutability in later chapters. For present purposes, it suffices to observe that the standard objections to Definability are not objec-

¹⁸ See Williamson 2000 for discussion of this point in the context of knowledge. Williamson 2007 suggests that common descriptions of Gettier cases do not suffice for the absence of knowledge, for example because there are deviant cases compatible with these descriptions in which subjects have other evidence for the relevant *p* (see Malmgren 2011 and Ichikawa and Jarvis 2009 for discussion). *G* above may escape this charge by including the “based solely on” clause. But the point still applies to justification: there will be deviant possible cases that satisfy *G* but not *J* because extraneous factors undermine Smith’s justification for believing the relevant proposition. Deviant cases undermine conclusive a priori scrutability (in the sense of 2.1) of *J* from *G* and may undermine any short nontrivial a priori sufficient condition for justification, but they do little to undermine the weaker scrutability claim that *J* is nontrivially a priori scrutable from a full enough specification of the case. An analogy: deviant cases undermine necessitation of *J* by *G* and may undermine any short nontrivial modally sufficient condition for justification, but they do little to undermine the weaker supervenience-style claim that *J* is nontrivially necessitated by a full enough specification of the case.

tions to A Priori Scrutability and that A Priori Scrutability remains an attractive thesis in the face of them.

4 From descriptions to intensions¹⁹

At this point we can take a leaf from Carnap's later work, especially his 1947 book *Meaning and Necessity*, and understand the meaning of expressions not in terms of definitions but in terms of *intensions*. Here the intuitive idea is that an intension captures the way an expression applies to possible cases of all sorts. For example, the Gettier case brings out that whether or not there is a good definition for 'know', we can classify different scenarios as involving knowledge or as not involving knowledge. An intension is a way to represent those classifications.

The intension of an expression can be identified with a function from scenarios to extensions, mirroring speakers' idealized judgments about the extension of the expression in the scenario. The intension of a sentence (as used in a context) is a function from scenarios to truth-values. For example, the intension of 'Smith knows that Jones owns a Ford or Brown is in Barcelona' is false in a Gettier scenario. The intension of a subsentential expression such as 'bachelor' is a function from scenarios to sets of individuals. In any given scenario, its intension picks out the people who are bachelors if that scenario is actual. An expression's intension will often depend on its context of use, but for simplicity I will set aside this context-dependence for now.

For our purposes, we can think of these scenarios as *epistemically possible* scenarios: roughly, highly specific ways the world might turn out that we cannot rule out a priori. (Here and throughout, I work with an idealized notion of epistemic possibility that is tied to what cannot be ruled out a priori.) For a given scenario w and a given sentence S , we can consider the hypothesis that w actually obtains and judge whether if w obtains, S is the case. If yes, the intension of S is true at w . If no, the intension of S is false at w . I give a fuller definition of scenarios and intensions in the tenth excursus, but for now we can work with this intuitive understanding.

On this model, speakers can grasp an expression's intension without grasping a corresponding definition. Instead, the grasp corresponds to a *conditional ability* to identify an expression's extension, given sufficient information about how the world turns out and sufficient reasoning. That is, a sufficiently rational

¹⁹ This section presupposes a little more philosophical background than the rest of the chapter and can be skipped without too much loss by nonspecialists. There is a somewhat gentler introduction to the framework of intensions in chapter 5, sections 3–5. A more precise account is in the tenth excursus.



subject using expressions such as ‘bachelor’, ‘knowledge’, and ‘water’ will have the ability to evaluate certain conditionals of the form ‘If E , then C ’, where E contains relevant information about the world (typically not involving the expression in question) and where C is a statement using the expression and saying whether a given case fall into its extension (e.g. ‘John is a bachelor’, ‘Sue knows that p ’, ‘Water is H_2O ’). And in order that it is not an accident that subjects can do this in the actual world, subjects will also be able to do this given specifications of many different scenarios.

For some expressions, we can capture the intension of the expression in the form of a definition. In other cases, we will merely be able to approximate an intension with an *approximate definition*. For example, ‘justified true belief’ can be seen as an approximate definition for ‘knowledge’: it gets most cases right, in an intuitive sense of ‘most’. ‘Justified true belief not essentially grounded in a falsehood’ is even better. In the face of counterexamples, one can refine definitions yielding longer and longer definitions that cover more and more cases. If there is no finite definition that gets all possible cases right, there may be a converging series of definitions: a series of longer and longer approximate definitions such that for any given case, there is some point in the series after which all definitions get that case right. In all these cases, however, the definitions are beholden to the intension rather than vice versa.

Arguments from counterexample can make a case against definitions, but they cannot make a case against the claim that expressions have intensions. Such arguments themselves proceed by considering scenarios (say, a Gettier scenario), and by making the case that the extension of an expression E (‘ S knows that P ’) with respect to that scenario differs from the extension of a purported definition D (‘ S has a justified true belief that P ’). To capture the intuitive data on the intensional model, we need only suppose that the intension of the expression picks out the intuitive extension at that scenario (in this case, false) rather than the intuitive extension of the definition (in this case, true).

All this applies equally to Kripke’s arguments against descriptivism, which are also arguments from counterexample. In fact, Kripke deploys two different sorts of arguments from counterexample. We might say that *modal* arguments from counterexamples are used to oppose the claim that ‘ N is the D ’ is necessary (for a name N and a description D), while *epistemic* arguments from counterexample are used to oppose the claim that ‘ N is the D ’ is a priori. In the case of knowledge, the Gettier counterexample serves as the basis of both a modal argument and an epistemic argument, showing that it is neither necessary nor a priori that knowledge is justified true belief. In Kripke’s arguments against descriptivism, modal arguments and epistemic arguments from counterexample are employed separately.

Modal arguments from counterexample require exhibiting a *metaphysically possible* situation (roughly, a situation that might have obtained) of which the equivalence is false. Kripke's modal argument against descriptivism fits this template. It focuses on a metaphysically possible situation in which Aristotle did not go into pedagogy, and makes the case that if this situation had obtained, then it would not have been the case that Aristotle was the teacher of Alexander. It follows that it is not necessary that Aristotle was the teacher of Alexander.

Epistemic arguments from counterexample require exhibiting an *epistemically possible* scenario (that is, a scenario not ruled out a priori) of which the equivalence is false. Kripke's epistemological argument against descriptivism is an argument from counterexample of this second kind. It focuses on an epistemically possible situation in which the proof of the incompleteness of arithmetic was stolen, and makes the case that if that situation actually obtains, then Gödel is not the prover of incompleteness. It follows that it is not a priori that Gödel is the prover of incompleteness.

In effect, modal arguments from counterexample show that the *modal profile* of an expression (the way it applies across metaphysically possible worlds) is not identical to that of a purported definition. Such an argument is clearly compatible with the claim that the modal profile can be represented as an intension, however. As usual, we need only choose an intension that respects the counterexample. The modal profile of 'know' can be represented as an intension that classifies Gettier cases as cases in which knowledge is absent. Likewise, the modal profile of 'Aristotle' can be represented as an intension that picks out Aristotle in the situation in which he never went into pedagogy, rather than picking out Alexander's teacher.

Similarly, epistemic arguments from counterexample show that the *epistemic profile* of an expression (the way that it applies across epistemically possible scenarios) is not identical to that of a purported definition. Again, such an argument is clearly compatible with the claim that the epistemic profile of an expression can be represented as an intension.²⁰ The epistemic profile of 'knows that P' can be represented as an intension that classifies Gettier cases as cases in which knowledge is absent. Likewise, the epistemic profile of 'Gödel' can be represented as an intension that picks out the stealer in Kripke's stolen-proof scenario rather than the prover.

²⁰ In the case of an expression such as 'knowledge', the epistemic and modal profiles appear to be more or less the same, so one intension will suffice to represent both. In the case of names such as 'Aristotle' and 'Gödel', the epistemic and modal profiles may be quite distinct, so one needs distinct intensions to represent them. These are just the primary and secondary intensions of two-dimensional semantics (discussed in 5.5 and Ero). The intension over epistemically possible scenarios discussed in the text is the primary intension, which is the most important for present purposes.

Like Gettier's argument from counterexample, Kripke's arguments from counterexample pose no problem for A Priori Scrutability. Kripke's modal argument does not pose even a *prima facie* problem: it concerns what is metaphysically possible and necessary, whereas A Priori Scrutability concerns what is a priori and a posteriori. Kripke's epistemological argument suggests that 'Gödel' is not a priori equivalent to a description such as 'the prover of incompleteness', but it gives no reason to deny that sentences such as 'Gödel did not prove incompleteness' are themselves scrutable from a specification of the relevant scenario. Given a specification of the stolen-proof scenario, we can certainly determine that if the scenario is actual, Gödel did not prove incompleteness.

Likewise, Kripke's epistemological argument cannot refute "approximate descriptivism": the thesis that for every name (as used by a speaker) there is a converging series of descriptions such that for every scenario, there is some point in the series such that all descriptions after that point give the same result as the name in that scenario. An approximate definition that works fairly well for 'Gödel' is 'The actual person called "Gödel" by those from whom I acquired the name'.²¹ As usual the approximation will be imperfect and there will be counterexamples (cases where one misheard the name, perhaps), but refinements will gradually remove the counterexamples as they converge on the name's intension. In any case, these counterexamples pose no more of a problem for A Priori Scrutability or for the intensional model than the Gettier case.

Much follows from these observations. Kripke's arguments are often thought to undermine broadly Fregean analyses of meaning and content. But we will see shortly (and in more detail in the eleventh excursus), an appropriate scrutability thesis can itself be used to support a broadly Fregean analysis of meaning and content, by defining intensions over epistemically possible scenarios. The resulting intensions can do much of the work that descriptions or Fregean senses are often held to do.

We can put things as follows. If the scrutability thesis is correct, a Fregean view of meaning and content is viable. Kripke's arguments give us no reason to reject the scrutability thesis. So Kripke's arguments should not lead us to reject a Fregean view of meaning and content. The scrutability thesis therefore suggests that Kripke's arguments are much more limited in scope than is often supposed. Of course there is more to say here, but this at least makes an initial case that the seemingly innocuous scrutability thesis may have highly significant consequences.

²¹ For more on approximate descriptivism, see 8.2. For more on intensions and approximate descriptions in the 'Gödel' case, see the discussion of Kripke's epistemological argument in "On Sense and Intension".

5 The scrutability base

A *scrutability base* is a class of truths from which all truths are scrutable, for a given notion of scrutability.²² What sort of truths might go into a scrutability base?

At the end of the *Aufbau*, Carnap embraced what we might call Logical Scrutability: the view that there is a scrutability base using only logical expressions. Some phenomenalists accept Phenomenal Scrutability, holding that there is a scrutability base using only phenomenal expressions (expressions for the character of conscious experiences) and logical expressions. Some physicalists accept Microphysical Scrutability, holding that there is a scrutability base using only microphysical expressions (expressions used in fundamental physics) and logical expressions.²³ For our purposes, all of these views are strong and interesting scrutability theses (versions of all of them are entertained by Carnap in the *Aufbau*), but the current project is not committed to any of them. Our working scrutability thesis is what we might call Compact Scrutability: there is a compact class of truths from which all truths are scrutable. Given that logical, microphysical, and phenomenal bases count as compact, then Logical, Phenomenal, and Microphysical Scrutability entail Compact Scrutability. But less austere bases than these may still be compact.

What is compactness, exactly? As I characterized compactness earlier, a class of truths is compact if it uses only a small class of expressions. A little more precisely, we can say that compactness requires that a class of truths uses only expressions from a small number of *families* of expressions. If it turns out that all truths are scrutable from phenomenal truths, but that an infinite number of phenomenal expressions are required to capture the diversity of possible phenomenal qualities, this would still be a strong enough scrutability thesis for our purposes. We can stipulate that the class of phenomenal expressions counts as a single family, as does the class of microphysical expressions, the class of logical expressions, the class of mathematical expressions, and so on. The intuitive idea here is that expressions in the same family should share a common domain. (So the class of spatiotemporal expressions counts as a family, while the class of singular terms does not.) Beyond this I will leave the notion of a family intuitive.

²² I will speak of sets and classes interchangeably. For some purposes it might be useful to admit classes of sentences that are too large to form a set, but for most of our purposes set-sized classes will be adequate. I discuss this issue further toward the end of E3.

²³ In principle, these views concerning a scrutability base can be combined with different scrutability relations (such as definitional or a priori scrutability), yielding such theses as Definitional Phenomenal Scrutability, A Priori Physical Scrutability, and so on. When the scrutability relation is not specified, a thesis involving a priori scrutability should be understood. For more on the conventions here, see 2.1.

How small is small? We can leave this notion vague. But to give a rough idea, I would say that fewer than ten or so families would be ideal, that twenty would be acceptable, but that more than a hundred would be pushing things. One could also stipulate that a compact class of truths will exclude the great majority of terms used in natural languages: there will be few or no ordinary proper names ('London', 'George Bush'), natural kind terms ('water', 'kangaroo'), artifact terms ('car', 'table'), and neither will there be cognate terms in a different language, constructions from such terms, and so on. The idea is that truths involving terms like this should all be scrutable from truths in a more primitive vocabulary. I will not build this into the official definition, but one can see this as part of the spirit of the thesis.

It is worth noting that while a compact class of truths must use only a limited vocabulary, it need not include *all* truths that use a given vocabulary. For example, there is a compact class of truths that includes all microphysical truths but not all mathematical truths. Stating the microphysical truths may require mathematical vocabulary, but many truths that use only mathematical vocabulary will not be included.

We also need to require that a compact class of truths avoids *trivializing mechanisms*. There are certain sorts of base truths that threaten to render the scrutability thesis trivial. One such is a base consisting of the family of expressions for *propositions*, along with 'is true'. It is not implausible that every sentence is scrutable from a sentence saying that a corresponding proposition is true, but this result is not interesting. Likewise, one could perhaps code all truths of English into a single real number φ , via an appropriate coding scheme: then it is not implausible that all such truths are scrutable from the single truth that φ equals such-and-such. But again, this thesis is not interesting. There is a clear sense in which these proposals involve trivializing mechanisms, by somehow directly coding a large number of truths from different families into a single truth or a single family of truths. I will not attempt to define this notion, but it should be understood that compact classes cannot include sentences of this sort.

So a class of sentences is compact if it includes expressions from only a small number of families and includes no trivializing mechanisms. Of course this notion is vague and has not been precisely defined. But in practice, this will not matter. The sort of specific scrutability claims I will discuss and defend will all involve highly restricted vocabularies that are clearly small enough to be interesting. In most cases, there will be no threat of a trivializing mechanism, and when there is such a threat, it can be discussed directly.

How small can a scrutability base be? Let us say that a *minimal* scrutability base is a class of sentences C such that C is a scrutability base and no proper subclass of C is a scrutability base. (In order to ensure that C uses a minimal vocabulary, one could also require that there is no scrutability base using only a proper subclass of the expressions used in C .)

Three proposals about minimal scrutability bases correspond to the theses of Logical Scrutability, Phenomenal Scrutability, and Microphysical Scrutability. I think that there are good reasons to reject these proposals, however. In part for reasons we have already discussed, it is plausible that many physical truths are not a priori scrutable from logical or phenomenal truths. Conversely, it is plausible that many phenomenal truths are not a priori scrutable from a microphysical base. For example, it appears that no amount of a priori reasoning from microphysical truths will settle what it is like to see red (Jackson 1982). This suggests that many phenomenal truths (truths concerning the character of conscious experiences) are not a priori scrutable from microphysical truths. It also appears that no amount of a priori reasoning from microphysical truths will enable one to know such perspectival truths as ‘It is now March’, or such negative truths as ‘There are no ghosts’.

Still, this leaves more liberal scrutability theses on the table. I will argue (in chapters 3, 4, and 6) that all ordinary macroscopic truths are a priori entailed by a class *PQTI* (physics, qualia, that’s-all, indexicals) that includes both truths of physics *and* phenomenal truths, as well as certain indexical truths (‘I am such-and such’, ‘Now is such-and-such’) and a “that’s-all” truth (on which more in 3.1 and E5). If so, then *PQTI* can serve as a scrutability base. There may be even smaller bases. For example, microphysical truths may themselves be scrutable from a base involving phenomenal expressions and nomic expressions (such as ‘law’ or ‘cause’), perhaps along with spatiotemporal and/or mathematical expressions. If so, then (as I argue in chapter 7) a scrutability base might need to involve only phenomenal, nomic, logical, indexical, and “that’s-all” expressions,²⁴ perhaps along with spatiotemporal and/or mathematical expressions. On some views (explored in chapters 7 and 8), the base may be smaller still.

A few principled scrutability bases are worthy of attention. One base, in the spirit of Carnap’s own view, yields the thesis of Structural Scrutability: all truths are scrutable from structural truths. If structural truths are restricted to a logical vocabulary, this view falls prey to Newman’s problem. But we might understand structural truths more expansively, to let in truths about fundamentality or naturalness (as on Carnap’s own final view), or about laws and causation, for example. I explore the viability of views of this sort in chapter 8.

Another principled scrutability thesis, perhaps less in the spirit of Carnap’s view, is Fundamental Scrutability: the thesis that all truths are scrutable from

²⁴ Throughout this book, I count as “indexical expressions” just a limited class of perspectival expressions: ‘I’, ‘now’, and perhaps certain heavily constrained demonstratives. In this sense, indexical expressions count reasonably as a family. I use “context-dependent” for the broader class of expressions whose content depends on context.

metaphysically fundamental truths (plus indexical truths and a that's-all truth, if necessary). The metaphysically fundamental truths are those that serve as the metaphysical grounds for all truths: they might involve attributions of fundamental properties to fundamental entities.²⁵ On a standard physicalist view, the metaphysically fundamental truths are microphysical truths. On a standard property dualist view, metaphysically fundamental truths may include microphysical and phenomenal truths.

Another thesis, in the spirit of Russell's quite different constructions of the world, is Acquaintance Scrutability: all truths are scrutable from truths about entities with which we are directly acquainted. Another, in the spirit of the thesis about concepts with which we started this chapter, is Primitive Scrutability: all truths are scrutable from truths involving only expressions for primitive concepts. Yet another, relevant to debates about internalism and externalism about meaning and content, is Narrow Scrutability: all truths are scrutable from truths whose content is determined by the internal state of the subject.

In chapter 8, I will make a case for all three of the theses just mentioned, as well as a tentative case for Fundamental Scrutability. I will also connect each of these theses to philosophical applications. For the purposes of many applications, it is these specific scrutability theses rather than Compact Scrutability per se that matters. Compactness plays a role in some applications, but where it does not, it can be seen as playing a sort of guiding role en route to the specific theses, ensuring that our scrutability bases are small enough that those theses are plausible.

Some potential scrutability bases are less austere than others. For example, someone might think that we need normative expressions ('ought') in the base, or that we need expressions for secondary qualities ('red') in the base, or that we need intentional notions ('believes') in the base. If a scrutability base needs to be expanded to include these expressions, then the base will plausibly go beyond the structural or the metaphysically fundamental, but it will still be small enough that we will have a strong and interesting scrutability thesis.

There are many scrutability bases. For a start, as long as scrutability is monotonic (if S is scrutable from C , S is scrutable from any set of truths containing C) adding truths to any scrutability base will yield a scrutability base, and substituting a priori equivalent synonyms within a scrutability base will also yield a scrutability base. Even if we restrict ourselves to minimal scrutability bases (scrutability

²⁵ Metaphysical fundamentality should be distinguished from conceptual primitiveness. One might reasonably hold that spin and charge are metaphysically fundamental without holding that the concepts *spin* and *charge* are primitive. Likewise, one might hold that the concept *I* is primitive without holding that the self is anything metaphysically fundamental. Still, there may be an attenuated relation between the two; see E16, and also 8.4 and 8.6.

bases of which no proper subclass is a scrutability base and for which there is no scrutability base using only a proper subset of the expressions) and factor out synonyms, a diversity of bases is possible. For example, given a minimal scrutability base involving predicates F and G , there will also be a minimal scrutability base involving four new predicates H, I, J , and K , corresponding to conjunctions of F, G , and their negations. One can also obtain multiple bases from the familiar idea that there can be a priori equivalent formulations of a physical theory in different vocabularies. It is even not out of the question that on some views, both a microphysical vocabulary and a phenomenal vocabulary (or a phenomenal vocabulary combined with a nomic or spatiotemporal vocabulary) could yield minimal scrutability bases.

For most of our purposes, the existence of multiple scrutability bases is not a problem. Carnap himself held a pluralistic view on which there are many equally privileged bases that we can choose between only on pragmatic grounds. Still, the phenomenon does suggest that the mere fact that an expression is involved in a minimal scrutability base does not suffice for the expression to express a primitive concept in an interesting sense. And there remains an intuition that some scrutability bases are more fundamental than others. For example, in the case above, it is natural to hold that predicates F and G stand in certain conceptual, epistemological, and psychological priority relations to H, I, J , and K . Likewise, one might hold that phenomenal and nomic expressions stand in certain conceptual, epistemological, and psychological priority relations to microphysical expressions. This will be especially clear if one holds that microphysical expressions are definable in terms of phenomenal and nomic expressions, but even if one rejects the definitional claim, one might still accept some priority claims.

I take the moral here to be that a priori scrutability is a relatively coarse-grained relation among classes of truths. One might react to this by postulating a more fine-grained relation of conceptual or epistemological dependence among truths. Whenever one class of truths depends on another in this sense, truths involving the former will be scrutable from truths involving the latter, but not vice versa. On this way of doing things, many scrutability bases will not be dependence bases, and it is not out of the question that there might be just one minimal dependence base (at least up to equivalence through synonymy).²⁶

²⁶ This reaction is an epistemological or conceptual analog of a familiar metaphysical line of thought concerning supervenience, leading some to postulate relations of ontological dependence or grounding that are finer-grained than the coarse-grained relation of supervenience. We could think of the more fine-grained relation as conceptual dependence or conceptual grounding. For more on these issues, see E16.

This line of thought immediately raises the question of how the fine-grained dependence relation in question should be understood. If one accepts the definitional model, one might suggest that the relation is just definitional scrutability, and that the dependence base will involve all and only the undefinable expressions. But if one rejects the definitional model, the correct understanding is less clear. I discuss such fine-grained relations and their relation to scrutability later in the book (and also in the companion chapter, “Verbal Disputes”).

For now, I will concentrate on a priori scrutability and related coarse-grained notions. These have the advantage of being better-understood than more fine-grained notions, so that arguing for scrutability theses of this sort is more straightforward. A number of the scrutability bases I will consider will also be plausible candidates to be dependence bases, so that the expressions involved will be plausible candidates to be primitive concepts. But even in the absence of claims about dependence and primitiveness, these scrutability theses have significant consequences.

6 Scrutability and the *Aufbau*

If the A Priori Scrutability thesis is correct, it offers a vindication of *something like* the project of the *Aufbau*.²⁷ There are two significant differences: the very limited bases (logical and/or phenomenal) of the *Aufbau* are replaced by somewhat less limited bases here, and the role of definitional entailment in the *Aufbau* is played by a priori entailment here. The expansion of the base allows us to avoid Goodman’s, Quine’s, and Chisholm’s objections to the phenomenalist base, and Newman’s objection to the purely logical base. The move from definitions to a priori entailment allows us to avoid the central problems for definitions and descriptions, including the problem of counterexample, and Kripke’s modal and epistemological arguments against descriptivism.

Of course there are challenges to the *Aufbau* that also apply to the scrutability framework. Most notably, Quine’s critique of the analytic/synthetic distinction is often thought to generate an equally significant critique of the a priori/a posteriori distinction, and so has the potential to undermine the A Priori Scrutability thesis. In chapter 5, however, I will suggest that an analysis in terms of scrutability provides the materials required to show where Quine’s arguments go

²⁷ A quite different project in a similar spirit, attempting to vindicate something like the *Aufbau*, is carried out by Hannes Leitgeb in “New Life for Carnap’s *Aufbau*?” (2011). Leitgeb retains a phenomenal basis, although he gives it more structure than Carnap allowed. He also retains definitional entailment by imposing a relatively weak criterion of adequacy according to which definitions must involve “sameness of empirical content”. On this criterion, definitions can be false. Because of this, I think that Leitgeb’s version of the *Aufbau* will not play the semantic, metaphysical, and epistemological roles that I am interested in, but it may well be able to play other roles.

wrong. I will address a number of other challenges to the scrutability framework in chapters 3 and 4.

One might ask: does A Priori Scrutability have the potential to satisfy some of the ambitions of the *Aufbau*? These ambitions included an analysis of meaning and concepts, an epistemological optimism, a metaphysical deflationism, and a language that might help to unify science. These elements were supposed to jointly yield a sort of blueprint for scientific analysis and philosophical progress. The *Aufbau* is widely held to have failed in these ambitions, and I will not try to put anything so strong in their place. Still, the scrutability thesis has consequences in many different areas of philosophy, consequences that share at least some of the flavor of Carnap's ambitions in the *Aufbau* and other works.²⁸

1. *Knowability and skepticism.* In the *Aufbau*, Carnap used his construction to argue that there is no question whose answer is in principle unattainable by science. This is a version of the notorious *Knowability Thesis* in epistemology, often associated with the programs of logical empiricism and verificationism, which holds that all truths are knowable. This thesis is now widely rejected, for both formal and intuitive reasons. I argue shortly (E1) that scrutability theses capture at least a plausible relative of these theses, and can play some parts of the role that the knowability thesis has been used to play. Furthermore, certain scrutability theses offer a distinctive response to skepticism (E15).

2. *Modality.* Carnap's *Aufbau* project yields a basic vocabulary that can be used not just to characterize the actual world, but also other possible states of the world. This leads directly to Carnap's later project in *Meaning and Necessity* (1947), in which he analyzes possibility and necessity in terms of state-descriptions for other possible worlds. While this sort of construction is now often used to understand metaphysically possible worlds, the scrutability framework allows such a construction to yield a space of epistemically possible worlds, or scenarios (E10). One can use a generalized scrutability thesis to define epistemically possible scenarios in terms of maximal a priori consistent sets of sentences in a scrutability base. These are analogous to Carnap's state-descriptions, and behave

²⁸ For more on these applications, see E1 and E15 (knowability and skepticism, respectively), E9 and E10 (modality and meaning), 8.3 and 8.4 (primitive concepts and narrow content), 8.6 and E16 (metaphysics), 8.7 and E12 (structuralism and the unity of science), and 6.5 (metaphilosophy). It should be noted that many of these applications require specific scrutability theses. For example, the reply to skepticism requires Structural Scrutability or a variant thereof. The analysis of narrow content requires Narrow Scrutability. Central applications to metaphysics require theses such as Fundamental Scrutability. The crucial applications to meaning and modality require less, but they work better if one at least has scrutability from a compact base consisting of non-context-dependent expressions and primitive indexicals, and better still if one has a version of Acquaintance Scrutability. See chapter 8 for a discussion of most of these matters.

in a more Carnapian way than possible worlds on the usual contemporary understanding. For example, a posteriori sentences such as ‘Hesperus is Phosphorus’ are true in all metaphysically possible worlds, but they are false in some epistemically possible scenarios, as one might expect. So these scenarios can play a role in analyzing epistemic possibility analogous to the role of possible worlds in analyzing metaphysical possibility.

3. *Meaning.* Carnap’s construction in *Meaning and Necessity* was intended to support a Fregean analysis of meaning, by understanding meanings as intensions defined over possible worlds. As discussed in chapter 5 and the eleventh excursus, the scrutability framework can be used to help vindicate this Fregean project by defining intensions over epistemically possible scenarios as above. For example, one can define the (epistemic or primary) intension of a sentence as the set of scenarios in which it is true. Then two sentences will have the same intension if and only if they are a priori equivalent. One can go on to define intensions for other expressions, such as singular terms, such that ‘a’ and ‘b’ will have the same intension if and only if ‘a = b’ is a priori. So ‘Hesperus’ and ‘Phosphorus’ will have different intensions. If the scrutability thesis is true, intensions of this sort will behave in a manner reminiscent of Fregean sense.

4. *Concepts and mental content.* In the *Aufbau*, Carnap put much emphasis on the construction of concepts. We can use the scrutability framework to associate intensions not just with linguistic items such as sentences but with mental items such as thoughts. As in the case of language, these intensions will serve as contents that reflect the epistemological properties of thoughts. Under some reasonable assumptions (outlined in the discussion of Narrow Scrutability in chapter 8), these intensions can also serve as *narrow* contents of thought: contents that are wholly determined by the intrinsic state of the thinker. These contents, grounded in a priori inferential relations to thoughts composed of primitive concepts, can go on to ground wide contents in turn. This approach to content naturally leads to a view in which primitive concepts play a grounding role with respect to all intentionality, and suggests that the path to naturalizing intentionality may proceed through the naturalization of the content of these primitive concepts.

5. *Metaphysics.* Carnap’s philosophy was known for its anti-realism about metaphysics: many metaphysical questions do not have objective and determinate answers. With specific scrutability theses in hand, the current framework can be used to argue for realism, anti-realism, or metaphysical primitivism about a given subject matter. For example, given Fundamental Scrutability (the thesis that all truths are scrutable from fundamental truths), then if ontological sentences (about the existence of composite objects, say) are not scrutable from more fundamental truths, then they are either themselves fundamental or they are not true. In the domain of ontology, one might use this method to argue for

a sort of anti-realism.²⁹ In other domains (that of consciousness, say), one can use this method to argue for an expansion in the metaphysically fundamental truths. We can also use scrutability as a guide in various projects of conceptual metaphysics, discussed in the sixteenth excursus.

6. *Scientific analysis.* The unity of science was one of the major concerns of the logical empiricists, and Carnap hoped that the *Aufbau* program might contribute to this unity by showing how all scientific notions could be analyzed in terms of a common basic vocabulary. If the scrutability thesis is true, then all scientific truths are at least scrutable from a common base. Furthermore, it can be argued that when scientific truths are scrutable from other truths of which there is a scientific account, this account can be used to provide an explanation of the scrutable truths. If so, then (as I argue in E12), scrutability might yield a relatively unified account of all scientific truths. Scrutability also helps to analyze the prospects for structuralist views of science (8.7).

7. *Metaphilosophy.* The scrutability thesis entails that all philosophical truths are scrutable from base truths. So even philosophical ignorance can be localized to our ignorance of base truths or the non-ideality of our a priori reasoning (6.5). An extension of the scrutability thesis (“Verbal Disputes”) suggests a way of reducing all philosophical disagreements to disagreements over base truths.

The analysis of meaning and concepts that one gets from this project is more open-ended than in the ambitions of the *Aufbau*, the epistemological optimism is attenuated, and any metaphysical deflationism is more limited. Still, the consequences are strong and striking enough that the scrutability thesis is certainly worthy of investigation.

²⁹ This application is restricted to distinctions between realism and anti-realism that can be drawn in terms of truth and falsity. The framework does not bear so directly on distinctions that are drawn differently: for example, arguments of this sort will not easily distinguish moral realism from varieties of moral anti-realism that allow that ‘Such-and-such is good’ is true. The framework itself is largely neutral on the nature of truth and its grounds in various domains. While I lean toward a correspondence view of truth myself, the arguments of this book are compatible with many different analyses of both realist and anti-realist flavors.

FIRST EXCURSUS

Scrutability and Knowability

The scrutability thesis is related to a number of widely discussed theses in analytic philosophy. In this excursus, I discuss the relation to the knowability thesis and its cousin the verification principle. In the next excursus, I discuss its relation to Quine's thesis of the inscrutability of reference. Doing so can help to indirectly motivate the scrutability thesis, by showing how it avoids problems for related theses while still capturing something of their flavor.

First, the Knowability Thesis.

Knowability Thesis: For any truth S , it is possible that someone knows S .¹

This thesis is often doubted, for both intuitive and formal reasons. Intuitively, it seems that there may be truths concerning the distant past, the far away, and the very small, that it may be impossible for anyone to know. Formally, the thesis gives rise to what is often known as the Paradox of Knowability, first published by Frederick Fitch in his 1963 article "A Logical Analysis of Some Value Concepts".²

Fitch in effect gives a disproof of the Knowability Thesis, arguing from the weak assumption that some truth is not *known* to the conclusion that some truth is not knowable. Let P be a truth such that in the actual course of history, no one ever knows P . Let Q be ' P and no one knows that P '. Then Q is true, but Q is unknowable. If someone were to know Q , then they would know P , but if someone were to know P , then Q would be false. So no one can know Q .

The scrutability thesis is closely related to the knowability thesis. It does not say that every truth is knowable, but it does say that every truth is scrutable, or derivable from a limited class of basic truths. One might thereby wonder whether

¹ I cast the thesis in terms of knowing sentences rather than knowing propositions for continuity with later discussion (see 2.2 for more on this). The present issues are much the same either way.

² Fitch attributes the result to an anonymous referee. Joe Salerno's "Knowability Noir: 1945–1963" locates the source in a 1945 referee report by Alonzo Church. The relevant material can be found in Salerno's *New Essays on the Knowability Paradox* (2009).

scrutability theses are liable to similar problems. In the introduction, we saw briefly that Inferential Scrutability is liable to problems related to Fitch's paradox, problems that I discuss at more length in 2.4 and 3.5. However, Conditional and A Priori Scrutability avoid both sorts of problems.

Concerning the intuitive problem: the base truths C may well include relevant truths about the distant past, including perhaps the spatiotemporal configuration of physical particles then, and so on. Even when S is an intuitively unknowable truth about the distant past, there is no corresponding intuitive problem with the idea that one can know that *if* the sentences in C are true, then S is true. Likewise there is no corresponding intuitive problem with the idea that one can know a priori a material conditional connecting a conjunction of all the truths in C to S . Something similar applies to truths about the far away and the very small. So there is no intuitive objection to the scrutability thesis here.

As for the paradox of knowability: even though Q above is unknowable, there is no formal problem with the claim that one can know that *if* the sentences in C are true, then Q is true. Indeed, as long as P itself and claims about knowledge of P are both scrutable from C , then ' P and no one knows that P ' will be straightforwardly scrutable from C . This goes for both A Priori and Conditional Scrutability.

One might suggest that the Scrutability Thesis entails the Knowability Thesis, at least if we grant that the conjunction of all truths in C is itself knowable. By knowing this conjunction D (empirically) and by knowing $D \rightarrow Q$ (a priori), one could thereby come to know Q . However, there is no reason to believe that D is itself knowable. In fact, there is good reason to believe that it is not, both for intuitive and Fitch-style reasons. The intuitive reasons are obvious: D may involve information about the distant past and the far away that no one will ever know. As for the Fitch-style reasons: assuming that no one in the actual history of the world believes D , then D specifies a world in which no one believes D . If someone came to believe D , they would live in a world quite different from ours, one in which their belief would be false. So no one can know D .

One might think that one can define a factive operator "scry" such that one scries P iff one derives P from base truths.³ One might then try to generate a Fitchian paradox, by taking P to be any truth that one does not actually scry, and taking Q to be ' P and I do not scry that P '. By Fitch's reasoning, if scrying is factive, then Q is an inscrutable truth. However: the notion of scrying above is ambiguous. If to scry P is to derive P from C , where C are the base truths of the actual world (or of any specific world) then scrying is not factive: there will be

³ "Scry" is the preferred verb form of "scrutable", having the advantage of both being more euphonious than the unlovely term "scrute" and already being a word of English with a somewhat appropriate meaning. "Scry: to divine, esp. by crystal gazing" (Collins English Dictionary).

worlds in which P is scried but false. If to scry P is to derive P from the base truths of the world one is in at the time of scrying, then ' P and I do not scry that P ' is indeed inscrutable. But this does not yield a counterexample to the A Priori or Conditional Scrutability theses above, as these require only that truths be derivable from the base truths of the worlds in which they are true, not the worlds in which they are so derived.

It may be that scrutability theses can do some of the work that knowability theses have been intended to do, or that they capture some of the intuitions that have led theorists to express sympathy with the knowability thesis. For example, Dorothy Edgington (1985) suggests that it is intuitive that if P is true in the actual world, then it is possible that one can come to know, in some different world, that P is true in the actual world. Of course this raises questions about what it is to know in a different world that P is true in the actual world. One suggestion is that to do this requires specifying the actual world with a canonical sentence D , and coming to know that if D were the case, P would be the case. Transposing this counterfactual claim into an epistemic mode (if D is the case, then P is the case), the resulting claim is not too far from the conditional scrutability thesis.

Another problem for the knowability thesis concerns cases of indeterminacy. (This problem is raised by Hawthorne (2005) for the case of omniscient knowers, but the problem generalizes.) Suppose that 42 is a borderline case of a small number, and let S be ' 42 is a big number'. On most views of vagueness, S is neither determinately true nor determinately false. On some such views, the statement $S \vee \neg S$ will be true all the same. If so, one could reason disjunctively: if S , then S is true, so S is knowable; if $\neg S$, then $\neg S$ is true, so $\neg S$ is knowable. So either S is knowable or $\neg S$ is knowable. But if S is (necessarily) indeterminate, this conclusion is implausible. One can raise a parallel problem for the scrutability thesis, yielding the conclusion that for all S , either S is scrutable (from a relevant D) or $\neg S$ is scrutable. Once again, this conclusion is implausible when S is indeterminate.

One could resist this conclusion by rejecting the law of the excluded middle and refusing to accept that $S \vee \neg S$ is true when S is indeterminate, or by holding that when S is indeterminate, it is likewise indeterminate whether S is scrutable. But perhaps the most straightforward way to avoid the problem is to understand the scrutability thesis as applying to *determinate* truths. That is, the thesis will say that when S is *determinately* true, or when $det(S)$ is true, then S is scrutable from D . On the relevant sort of view, the disjunction $det(S) \vee det(\neg S)$ will not be true in cases of indeterminacy, so the problem here will be avoided.

One might worry about cases of higher-order indeterminacy, where it is indeterminate whether $det(S)$ or $\neg det(S)$. In such a case, the best thing to say is that it is indeterminate whether S is scrutable from D . Given the presence of vagueness

in language, one should expect that scrutability can be vague too. On this view, implication by base truths goes along with determinacy, and vagueness of implication goes along with vagueness of determinacy.⁴

One can extend the scrutability thesis to the thesis that for all S , the truth-value of S is scrutable from D , whatever this truth-value may be. To obtain the extended thesis, one could simply apply the original thesis to the statement ‘ S has truth-value T ’, or better, one could apply the thesis to a statement such as ‘ $\neg S$ ’, ‘ $\text{indet}(S)$ ’, and other statements which are true iff S has a relevant truth-value. As in the cases above, then if one adopts the view of indeterminacy outlined above, these statements will be scrutable only when they are determinately true. So, for example, the claim will be that if $\text{indet}(S)$ is determinately true, then it is scrutable from D .

A final worry related to these matters arises from cases analogous to the Liar Paradox. Say that S is ‘This sentence is not scrutable from D ’. Then if S is true, it is inscrutable, and if S is false, it is scrutable. Either way we have a counterexample to the thesis that a sentence is true if and only if it is scrutable.

This worry is an instance of a general worry for any thesis holding that a sentence is true iff it has property φ . Whether ‘This sentence does not have φ ’ is true or false, it generates a counterexample to the thesis. I do not think it is reasonable to infer that no such thesis can be true. If this were correct, the Liar Paradox would generate a counterexample to ‘Every sentence is true iff it is true’. Instead, it seems best to say that sentences like ‘This sentence does not have φ ’ should be handled by whatever mechanism best handles the Liar Paradox. Indeed, one might take it to be a constraint on solutions to the Liar Paradox that they should also apply to sentences like this.

The most obvious thing to say is that in cases like this, ‘ S does not have φ ’ is indeterminate. Given the discussion above, ‘This sentence is not scrutable from D ’ is slightly more complicated, as the relevant thesis says that a sentence is *determinately* true iff it is scrutable. This renders the sentence at issue more closely analogous to the Strengthened Liar, ‘This sentence is not determinately true’. So a proponent of the Scrutability Thesis should say that the sentence has the same truth-value of the Strengthened Liar, whatever that truth-value is

⁴ This view is analogous to the view that knowability goes along with determinacy, and vagueness of knowability goes along with vagueness of determinacy, suggested on behalf of the supervaluationist by Hawthorne 2005. There is an alternative view (Dorr 2003) on which the vagueness of knowability goes along with the vagueness of truth rather than the vagueness of indeterminacy. Transposed to the key of scrutability, this approach yields a view on which S is true iff S is scrutable and \bar{S} is indeterminate iff it is indeterminate whether S is scrutable. If we accept the law of the excluded middle, this view will most naturally be combined with a view on which it is always the case that either S is scrutable or \bar{S} is scrutable (cases apparently in the middle will be borderline cases of each).

(perhaps involving some sort of higher-order indeterminacy). Saying more requires an adequate treatment of Liar paradoxes in general, but that is a problem for everyone, and not for the scrutability thesis in particular.

Finally, the scrutability thesis is in some limited respects reminiscent of the logical empiricists' verification principle, which says that only verifiable statements are meaningful. The scrutability thesis, rephrased, says that only scrutable statements are true, where a statement is scrutable if it is implied by certain base statements. Perhaps scrutability here might be seen as a sort of idealized verifiability, conditional on those statements in the base. One might then wonder whether any of the famous problems for the verification principle will apply here.

Most traditional worries about verifiability are removed by the extension of the base. Scrutability is much weaker than verifiability, not least because the base statements may include truths that are not themselves verifiable. For example, they may include truths about the distant past, the far away, about other minds, and about the extent of the universe. Because of this, there is no problem for scrutability generated by distinct empirically equivalent theories in physics, for example, or by statements about the past, or by the possibility of unverifiable ghosts.

Another famous problem is: is the verification principle itself verifiable? One might likewise ask: is the scrutability thesis itself scrutable? I will argue later that certain general versions of the scrutability thesis are themselves a priori, and are therefore scrutable. Other versions, such as scrutability from a specific base, are a posteriori. But we will later see that as long as a sort of "that's-all" sentence is included in the base, the scrutability thesis itself will follow. In some cases this "that's-all" sentence will itself be akin to a scrutability thesis, but this just brings out a way in which the scrutability thesis is far more flexible than the verification principle.

It is also worth noting that where the logical empiricists offered the verification principle in a prescriptive spirit, I am not inclined to offer the scrutability thesis in this way. Instead, in the first instance I am simply arguing for its truth. Perhaps downstream from these arguments, it can be used prescriptively, as a check on realism about certain subject matters that are not scrutable from base truths. Much here will depend on what one antecedently allows into the base, so the matter is not cut and dried. But in any case, it seems clear that the standard reasons for doubt about the verification principle do not apply to the scrutability thesis.