

9-2017


# Scientifically Responsible Metaphysics: A Program for the Naturalization of Metaphysics

Amanda Bryant

*The Graduate Center, City University of New York*

## How does access to this work benefit you? Let us know!

Follow this and additional works at: [https://academicworks.cuny.edu/gc\\_etds](https://academicworks.cuny.edu/gc_etds)

 Part of the [Metaphysics Commons](#), and the [Philosophy of Science Commons](#)

---

### Recommended Citation

Bryant, Amanda, "Scientifically Responsible Metaphysics: A Program for the Naturalization of Metaphysics" (2017). *CUNY Academic Works*.

[https://academicworks.cuny.edu/gc\\_etds/2244](https://academicworks.cuny.edu/gc_etds/2244)

This Dissertation is brought to you by CUNY Academic Works. It has been accepted for inclusion in All Dissertations, Theses, and Capstone Projects by an authorized administrator of CUNY Academic Works. For more information, please contact [deposit@gc.cuny.edu](mailto:deposit@gc.cuny.edu).

**SCIENTIFICALLY RESPONSIBLE METAPHYSICS:  
A PROGRAM FOR THE NATURALIZATION OF METAPHYSICS**

by

AMANDA BRYANT

A dissertation submitted to the Graduate Faculty in Philosophy in partial fulfillment of the requirements for the degree of Doctor of Philosophy, The City University of New York

2017

© 2017

AMANDA BRYANT

All Rights Reserved

Scientifically Responsible Metaphysics:  
A Program for the Naturalization of Metaphysics

by

Amanda Bryant

This manuscript has been read and accepted for the Graduate Faculty in Philosophy  
in satisfaction of the dissertation requirement for the degree of Doctor of  
Philosophy.

---

Date

---

David Papineau  
Chair of Examining Committee

---

Date

---

Iakovos Vasiliou  
Executive Officer

Supervisory Committee:

Michael Devitt

David Papineau

Peter Godfrey-Smith

THE CITY UNIVERSITY OF NEW YORK

## ABSTRACT

Scientifically Responsible Metaphysics: A Program for the Naturalization of Metaphysics

by

Amanda Bryant

Advisor: Michael Devitt

There has been much recent work calling for the naturalization of metaphysics, including most famously James Ladyman and Don Ross' polemic, *Every Thing Must Go*. But much work remains to adequately articulate and motivate the call to naturalize metaphysics. My dissertation contributes to that work. Its central questions are: *What relationship should metaphysics have to current science? Must good metaphysics be responsive to current science, and if so, how? Why should metaphysics be naturalized and what should its naturalization consist in?*

I argue, first, that for that for epistemic purposes, as opposed to heuristic or pragmatic purposes, theories should be robustly constrained and adequately warranted. The negative portion of the dissertation attacks what I call *free range metaphysics* — metaphysics, the content of which is constrained not by science, but only by logical, aesthetic, and psychological demands, such as the demands for consistency, simplicity, intuitive plausibility, and explanatory power. I argue that, individually and jointly, the constraints on the content of free range metaphysics are insufficiently robust and their satisfaction fails to secure sufficient epistemic warrant. Therefore, free range metaphysics cannot claim to produce justified theories of reality. The positive portion of the dissertation prescribes *scientifically responsible metaphysics* — metaphysics conscientiously engaged with the theories and practices of the current sciences. I argue that scientifically

responsible metaphysics is better constrained and supported than free range metaphysics and therefore can better claim to justify metaphysical theories.

Finally, I consider potential problem cases — cases in which some metaphysical topic is not *obviously* apt for being made scientifically responsible — including modal metaphysics and grounding. I resolve the problem cases by showing how the topics can be made scientifically responsible. First, I articulate a methodology for scientifically responsible modal metaphysics that takes current science as an evidence-base for the justification of modal claims and as a model of good modal reasoning. Second, I synthesize a list of fruitful uses of science for grounding theorists, including among other things: to help in the identification of putative grounding relata, to show correlations among them, to demonstrate their non-identity, to provide a stock of explanatory patterns, to identify candidate essential properties, and to motivate agnosticism about particular grounding theses where scientific support is lacking. Having resolved the problem cases, I conclude that the prospects for making metaphysics scientifically responsible are bright.

## ACKNOWLEDGMENTS

Very special thanks to my supervisor, Michael Devitt, for his patient dedication, careful and illuminating feedback, advice and support, motivation and encouragement, and good humour.

For their invaluable feedback and encouragement, I am also indebted to Graham Priest, David Papineau, Peter Godfrey-Smith, Jesse Prinz, Gary Ostertag, Barbara Montero, and Stephen Neale.

Warm thanks to my colleagues and friends Yuval Abrams, Cosim Sayid, Derek Skillings, Ralph Jenkins, Leonard Finkelman, Dustin Olson, and Paul Simard-Smith for illuminating conversations and written exchanges pertaining to this work. Thanks also to the members of Priestclub who listened to and commented on presentations of this work, and to the participants of Stephen Neale's Fall 2015 dissertation seminar, including especially Jake Quilty-Dunn and Ross Colebrook.

For helpful questions and comments, many thanks to audience members at my talks at the following conferences: *Canadian Philosophical Association Annual Congress 2015*, May 31, 2015 at the University of Ottawa; *Directions in the Epistemology of Modality*, October 22, 2015 at the University of Stirling; *New Trends in the Metaphysics of Science*, December 16, 2015 at the Université Paris-Sorbonne; *The Metaphysics and Epistemology of Grounding and Fundamentality Conference*, December 11, 2015 at the CUNY Graduate Center; and *Canadian Philosophical Association Annual Congress 2016*, May 29, 2016 at the University of Calgary. I am especially grateful to my commentators, including Brandon Beasley, Kyle Blanchette, and Christian Nitz.

This research was supported by the Social Sciences and Humanities Research Council of Canada, award #752-2012-0363. I thank them for their support, which has aided me significantly in my graduate career.

Finally, my heartfelt thanks to my family for their unwavering love and support – especially to Mom and Dad.



## TABLE OF CONTENTS

Abstract	iv.
Acknowledgements	vi.
<b>Chapter 1: The Call to Naturalize Metaphysics</b>	1
I. The Issue	2
II. Naturalism: A Brief History	6
III. <i>Physics-Venerating Verificationist Unificationism</i> Reconsidered	16
III.1 The Need for a Better Argument for Naturalization	16
III.2 The Need for a Better Program of Naturalization	18
III.2.1 Verificationism	19
III.2.2 Modes of Engagement	21
III.2.3 Unificationism	22
III.2.4 Physics Primacy	23
III.3 Desiderata for the Call to Naturalize Metaphysics	24
IV. Armchair Metametaphysics?: Methodological Considerations	25
V. Plan for the Dissertation	31
<b>Chapter 2: Epistemic Principles</b>	34
I. Epistemic Normativity	34
II. Epistemic Principles	41
II.1 Theoretical Constraint	41
II.2 Epistemic Warrant	48
III. Defending the Epistemic Principles	51

III.1 Constraint Principles	51
III.1.1 Robust Constraint and Simplicity	61
III.2 Warrant Principles	62
<b>Chapter 3: On the Inadequacy of Free Range Metaphysics and the Appeal of Scientifically Responsible Metaphysics</b>	<b>68</b>
Part I: Free Range Metaphysics	69
I.1 Definition	69
I.2 Examples	69
I.3 Failure to Meet Epistemic Criteria	72
I.3.1 Consistency	73
I.3.2 Simplicity	74
I.3.3 Intuitive Plausibility	76
I.3.4 Explanatory Power	81
I.3.5 No Joint Sufficiency	87
I.4 Collateral Benefits	88
I.4.1 Clarity	88
I.4.2 Fruitfulness in Science	90
I.4.3 Fruitfulness in Scientifically Engaged Philosophy	91
I.5 Normative Conclusion	93
Part II: Scientifically Responsible Metaphysics	97
II.1 Definition	97
II.2 Clarifications	98
II.3 Examples	101

II.4 Satisfaction of Epistemic Criteria	102
II.5 Normative Conclusion	104
Part III: Objections and Replies	104
1. Why go beyond science?	104
2. What role for philosophers?	105
3. Objection from underdetermination	106
4. Objection from scientific anti-realism	108
5. Objection from new cynicism	108
6. Objection from relevance	110
<b>Chapter 4: Case Studies</b>	112
Case I: Modal Metaphysics	113
I. The State of Modal Metaphysics, Epistemology and Practice	114
I.1 Modal Rationalism	114
I.1.1 Conceivability	115
I.1.2 Understanding	116
I.1.3 Further Difficulties	118
I.2 Williamson's Counterfactual Account	120
I.3 Modal Empiricism	122
II. Scientifically Responsible Modal Metaphysics	125
II.1 Science and Modal Truth	126
II.2 'Further Difficulties' Resolved	129
III. Conclusion, Case I	131
Case II: Grounding	131

I. The State of the Metaphysics of Grounding, Epistemology and Practice	133
I.1 Epistemology	133
I.1.1 Fine	133
I.1.2 Schaffer	135
I.1.3 Audi	137
I.2 Practice	141
I.2.1 Levels	142
I.2.2 A Fundamental Level	144
I.2.3 What Grounds What	147
I.2.4 Example-Based Method	150
II. Scientifically Responsible Metaphysics of Ground	153
II.1 Hofweber's Challenge	157
III. Conclusion, Case II	159
<b>Conclusion</b>	161
<b>Bibliography</b>	163

## CHAPTER 1

### The Call to Naturalize Metaphysics

Analytic metaphysics is bunk. Or so its critics say. A naturalistic impulse has gripped much recent philosophy and taken in its critical sights our more esoteric speculative inquiries, including much contemporary analytic metaphysics. The call to naturalize metaphysics is the call for it to make contact with science. Its central motivations are a distrust of philosophical intuition and other *a priori* methods, respect for the epistemic distinction of science, and the belief that there are fruitful points of contact between science and metaphysics. While Quine initiated the call to naturalize metaphysics in the mid-twentieth century, the prospect of naturalization has received renewed interest since Ladyman and Ross' infamous (2007) polemic, *Every Thing Must Go*. But there is much work to be done to better motivate and refine the call to naturalize metaphysics. This dissertation aims to contribute to that work.

In this first chapter, I will situate my project in its broader historical context. For brevity, the story I tell will be simplified and incomplete. I will trace the history of the call to naturalize metaphysics, from early modern empiricism, logical positivism, Quine, and the rise of analytic metaphysics, to Ladyman and Ross' seminal (2007) book, its reception, and developments since. I will also speak to the value of the project and the need for further refining and motivating the call to naturalize metaphysics. In particular, I will point to the need to clearly delineate science from non-naturalized metaphysics, to avoid the historical failures of the positivist program, to avoid a needlessly narrow conception of the work of naturalized metaphysics, to recognize the complex and pluralistic nature of the current sciences, and to not insist on the primacy of physics over and above other sciences. Lastly, I will describe and defend the methodology of the dissertation. In

particular, I will consider whether the central questions of the dissertation, like the central questions of metaphysics, should be settled by appeal to science. I will argue that first-order and second-order metaphysics do not stand or fall together, since their domains differ substantially. First-order metaphysics is about the world; second-order metaphysics is about metaphysics. Moreover, the sort of second-order metaphysics I do in this dissertation is *normative*, unlike first-order metaphysics. And, I argue, normative second-order inquiries must bottom out in some *a priori* evaluative criteria. Moreover, I am dubious of the claim that we could scientifically investigate the truth-conduciveness of different metaphysical methods. So my conclusions in this dissertation will not undermine the *a priori* arguments that establish them. Finally, to conclude the chapter, I will outline my plan for the rest of the dissertation.

## I. The Issue

The central questions at issue in the dissertation will be: *What is the proper relationship between metaphysics and current science? Must good metaphysics be responsive to current science, and if so, how?* To situate the issue with respect to recent literature, the question may be framed: *Why should metaphysics be naturalized and what should its naturalization look like?*

For metaphysics to be *naturalized* is for it to bear some appropriate relation to science. The relation may be characterized in a number of ways: naturalized metaphysics is *motivated* by, *responsive to*, *engaged with*, *answerable* to science, and so on. Any of these notions will do, but each of them needs spelling out. Let's stick with the notion of *engagement*.<sup>1</sup> Different advocates of naturalization prescribe different modes. For instance, Ladyman and Ross (2007) think that the right mode of engagement is *unifying scientific theses*. Maclaurin and Dyke (2012) think that the

---

<sup>1</sup> Not much hangs on word choice here. The importance lies in how the details are fleshed out.

right mode is *premissing metaphysical claims on scientific results*. In the picture of naturalized metaphysics — which I will call *scientifically responsible metaphysics* — I develop in Chapter 3, I will recommend a broad array of modes. What unites advocates of naturalization is their shared view that wild metaphysical speculation should be brought down to earth and that the products and practices of science have some crucial role in doing so.

The call to naturalize is often motivated by a commitment to naturalism, of both the epistemological and methodological variety.<sup>2</sup> Quine (1969, 1981) famously defended *epistemological naturalism*. The view entirely rejects the possibility of *a priori* knowledge and justification, and views scientific evidence as the only legitimate kind of evidence. Epistemological naturalism implies *methodological naturalism*, which requires that philosophical inquiry proceed by appeal to science and via scientific methods. Though epistemological and methodological naturalism can motivate the call to naturalize metaphysics, the call does not imply any commitment to these naturalisms. I will avow neither, since I do not wish to deny all *a priori* knowledge, nor do I wish to prescribe the naturalization of *all* domains. I will motivate my methodological recommendations differently.

Now, to get clear on the central questions of the dissertation, we must first have some clear conception of metaphysics. But I must define *metaphysics* carefully, to ensure that the notion is acceptable to the metaphysician and that it does not beg the question in my own favor. But it is no easy task to define metaphysics, much less to define it uncontroversially and neutrally. As Karen Bennett (2016) shows, many popular definitions of metaphysics are inadequate. According to Bennett, a definition of metaphysics must do three things. It must: 1) somewhat distinguish

---

<sup>2</sup>*Ontological naturalism*, according to which the world contains nothing supernatural (Papineau 2014), is not directly relevant.

metaphysics from science, 2) give metaphysics a subject-matter even if there is no fundamental level, and 3) reflect the actual practices of metaphysicians (2016, 29-30). Definitions according to which metaphysics attempts to “‘get behind the appearances’ and study ‘what the world is like’” fail to respect 1) (2016, 29). The view that metaphysics studies the fundamental nature or structure of reality fails to respect 1), 2), and 3) (2016, 29-30). The Aristotelean conception of metaphysics as the study of being *qua* being fails 3), since metaphysicians investigate other categories and particular phenomena (2016, 31). So in Bennett’s view, several popular definitions of metaphysics don’t pass muster.

I agree with Bennett, but I suggest a revision to her second criterion. In my view, the problem with definitions according to which metaphysics studies *fundamentalia* is *not* that they rob the discipline of a subject-matter if there are no *fundamentalia*, but that they beg the question against legitimately *metaphysical* views according to which there are none. So the second criterion should require that we don’t beg the question in that manner. In addition to ruling out the *study of the fundamental* definition, this revised criterion would rule out definitions claiming metaphysics to be the study of *objective* reality, which may beg the question against some (though not all) idealists.<sup>3</sup>

These problems are reflective of the terrible difficulty of delineating disciplines in a principled manner. Because of that difficulty, I can only gesture toward the sort of thing metaphysics does. Metaphysics describes aspects of reality at a relatively high level of abstraction.<sup>4</sup> In Craig Callender’s words, metaphysics makes claims about the world “more abstract and distantly related to experiment than... science” (2011, 47). Specifically, it concerns what exists

---

<sup>3</sup> I thank Jesse Prinz for this example.

<sup>4</sup> I thank Jesse Prinz for suggesting this gloss.



(ontology), the nature of the existents, the relations that obtain among them, as well as the nature of the properties and relations. For instance, it asks whether numbers, minds, substances, mereological sums, and universals exist fundamentally. It asks about the nature of causation, the conditions for object identity and persistence, and so on. The only definition I will give is an ostensive one — *metaphysics* is whatever it is that we do in metaphysics anthologies, journal articles, and classrooms.

I must also define *current science*. Like the concept *metaphysics*, the concept *science* is notoriously difficult to define, as our unresolved struggle with the demarcation problem shows. I will dodge the difficulty by again giving an ostensive definition. I define *current science* (sometimes just *science*) by pointing to the institution of science as embodied in laboratories and university departments around the world. It includes the practices, experimental data, and theoretical products of the people working in those contexts. I should note that I do not wish to limit my conception of current science to physics. I mean science broadly and pluralistically construed — natural and social — defined institutionally. This dissertation is about the proper relationship of metaphysics to current science, so understood.

Let me give some points of clarification. First, the topic of the dissertation is *normative*, not descriptive. I am concerned not with the relationship that the discipline of metaphysics *actually* stands in to current science, but the relationship that it *ought* to stand in, given the primary goal of metaphysical inquiry: knowledge of metaphysical fact, or more humbly, justified theories about the metaphysical facts. Further, the topic is both epistemological and methodological. It is epistemological insofar as it asks about the conditions for justification with regard to metaphysical truth. It is methodological insofar as it asks about the conditions for responsible metaphysical practice.

Second, although it is interesting and important to think about what science can, or should, or *does* take from metaphysics, I bracket that question for the purposes of this dissertation. I do think that certain metaphysical assumptions underlie scientific theorizing. So I do not wish to recommend some caricature view according to which science, being pure of metaphysics, comes *first* and metaphysics *second*. Rather, I think there ought to be a dialectic between them — a complementary, ongoing, back-and-forth relationship. Nevertheless, for my purposes, I am concerned with just the one side of that relationship.

## II. Naturalism: A Brief History

The desire to make metaphysics responsive to what we now call *natural science* has its historical roots in early modern empiricism, most famously exemplified by Locke, Berkeley, and Hume. Empiricism developed in opposition to the rationalism exemplified by Descartes, Spinoza, and Leibniz, who developed metaphysical systems largely through *a priori* speculation.<sup>5</sup> While the rationalists proceeded roughly independently of empirical evidence, empiricists considered it central. Where empiricism was not followed to its skeptical conclusion, it was friendlier to science than rationalism. Though it did not *mandate* the use of science, it privileged the broader class of empirical evidence that encompasses scientific evidence. In that respect, empiricism is an important part of the history of naturalism.

In the first half of the twentieth century, the logical positivist movement initiated by the Vienna Circle took empiricism a step further. The positivists accepted the empiricists' claim that empirical evidence is a requisite of knowledge. But they added to it an explicit deference for

---

<sup>5</sup> By tracing the history of naturalism back to empiricism, I do not mean to suggest that science can be equated with *a posteriori* methods. I only mean to suggest that respect for science emerges naturally from the methodological view that empirical evidence should be privileged.

science. The movement opposed “scholastic metaphysics” and in particular, “the systems of German idealism” inspired by Kant and Hegel (Vienna Circle 1929, §2). Though Kant took inspiration from Hume (Kant 1785/2001, 5) and rejected the possibility of gaining metaphysical knowledge from pure reason, he nonetheless thought that pure reason yields some substantive knowledge. Against this, the positivists held that “[t]he scientific world-conception knows no unconditionally valid knowledge derived from pure reason, no 'synthetic judgments a priori' of the kind that lie at the basis of Kantian epistemology” (Vienna Circle 1929, §2). They also rejected what is sometimes called *Hegelian metaphysics*. The attribution of a speculative metaphysics to Hegel is now controversial (Stern 2009). However, the logical positivists reacted against an approach to metaphysics commonly attributed to Hegel, according to which metaphysics is:

an attempt to characterize those supersensible entities or unconditioned objects that form its basic subject-matter, such as God, the soul, and the world as it came to be in space and time. Hegel, it seems, sees in our concern with such entities a desire to get beyond the ordinary, empirical world, and to cognize ‘the absolute’. (Stern 2009, 30)

On the approach commonly attributed to Hegel, metaphysics proceeds by logical investigation of such non-empirical concepts. The positivists’ reacted to the approach as follows:

If someone asserts ‘there is a God’, ‘the primary basis of the world is the unconscious’, ‘there is an entelechy which is the leading principle in the living organism’; we do not say to him: ‘what you say is false’; but we ask him: ‘what do you mean by these statements?’ (Vienna Circle 1929, §2)

As the final sentence suggests, the logical positivists’ critique of metaphysics was largely motivated by their theory of meaning. The logical positivists famously accepted a verificationist criterion of meaning, according to which sentences are only meaningful if they are empirically

verifiable. They thought metaphysical claims failed the criterion: “the statements of the old metaphysics... are meaningless, because unverifiable and without content” (Vienna Circle 1929, §2). In particular, Carnap (1950) denied the meaningfulness of metaphysics on the grounds that it attempts to ask framework-external questions. Ultimately, logical positivism failed, in part due to its verificationism:

The technical aims of positivism and logical empiricism—to show how all meaningful discourse can be reduced to, or at least rigorously justified by reference to, reports of observations regimented for communication and inference by formal linguistic conventions—have shown to be unachievable. (Ladyman and Ross 2007, 8).

Notwithstanding the failure of logical positivism, from the naturalists’ perspective, it had a noble spirit: it added to the empiricist stance an explicit respect for science and demanded of philosophical methodology that it privilege scientific evidence.

In the wake of the failure of logical positivism, Quine (1969, 1981) defended the possibility of metaphysics, while preserving the positivists’ respect for science. His rejection of the analytic-synthetic distinction amounted to a rejection of Carnap’s criticism of metaphysics: no analytic-synthetic distinction, no frameworks; “[n]o frameworks, no framework pluralism — and no place to banish metaphysics” (Callender 2011, 35). Further, Quine articulated and defended epistemological naturalism. In his words, epistemological naturalism is the “abandonment of the goal of a first philosophy prior to natural science” (1981, 67). More positively, it is, “the recognition that it is within science itself, and not in some prior philosophy, that reality is to be identified and described” (1981, 21). Such naturalism, he says, “sees natural science as an inquiry into reality, fallible and corrigible but not answerable to any supra-scientific tribunal” (1981, 72). Quine’s articulation of epistemological naturalism initiated the ongoing trend in the philosophy of

science toward epistemological and methodological naturalism (Devitt 1996, 1998; Giere 1988; Kitcher 1992; Laudan 1990c; Maddy 2007; and Papineau 2014). Importantly, Quine applied his epistemological naturalism to metaphysics and in so doing, re-popularized belief in the possibility of, and even the potential respectability of, metaphysics. In his view, naturalizing metaphysics requires that:

we adopt... the simplest conceptual scheme into which the disordered fragments of raw experience can be fitted and arranged. Our ontology is determined once we have fixed upon the over-all conceptual scheme which is to accommodate science in the broadest sense. (1951, 16-17)

We settle our ontology by appeal to our best science, the system into which it figures, plus considerations of simplicity. Quine helped resurrect metaphysics from its apparent death at the hands of the positivists, but he reined in its wilder speculative character with a naturalistic methodology.

Following Quine's vindication and methodological revision of the project of metaphysics, analytic metaphysics came to prominence in the latter half of the twentieth century. *Analytic metaphysics* is the movement popularized by Lewis and Armstrong that aims to apply analytic rigour to metaphysical topics like causation, universals, and modality. While analytic metaphysics carved out its subject-matter, Kripke convinced philosophers that common sense intuition has an important place in philosophy. He suggested, in Callender's words, that our robust intuitions "carve out a realm of modality not obviously reducible to logical or scientific possibility, namely metaphysical modality" (2011, 36). And, as Ladyman and Ross put it, ultimately, "analytical philosophy broke free of its Quinean moorings" (2007, 9 fn. 8). We eventually saw:

the resurgence of the kind of metaphysics that floats entirely free of science. Initially granting themselves permission to do a bit of metaphysics that seemed closely tied to, perhaps even important to, the success of the scientific project, increasing numbers of philosophers lost their positivistic spirit. (Ladyman and Ross 2007, 9)

Inspired by Quine, philosophers accepted the possibility of metaphysics, but they largely abandoned the naturalistic methodology on which its credibility depended. The result, as Ladyman and Ross characterize it, is that contemporary metaphysics is now full of “esoteric debates about substance, universals, identity, time, properties, and so on, which make little or no reference to science, and worse, which seem to presuppose that science must be irrelevant to their resolution” (2007, 10). Much analytic metaphysics now violates Quine’s call to naturalize.

In the early aughts, experimental philosophy emerged as a critical reaction against the reliance of contemporary philosophy on armchair intuitions. Proponents of experimental philosophy (Knobe 2007; Knobe and Nichols 2007; Prinz 2008; Sytsma and Livengood 2015; Sytsma and Machery 2009; Weinberg, Nichols and Stich 2001) attempt to import empirical methods, typically those of psychology, into the treatment of philosophical questions. They use those methods to study folk intuitions about language, ethics, mind, and epistemology, among other things. Some experimental philosophers have done studies pertaining to metaphysics, in which they study the ordinary understanding of, e.g. causation (Alicke, Rose, and Bloom 2011; Hitchcock and Knobe 2009; Livengood and Machery 2007), personal identity (Bruno and Nichols 2010) and mereology (Rose and Schaffer 2015). Critical of the non-experimental methods of contemporary metaphysics and of philosophy more broadly, experimental philosophers consign the armchair to the flames.

Following the inception of experimental philosophy in the early aughts, in 2007, Ladyman and Ross issued a scathing condemnation of speculative, “Neo-Scholastic” metaphysics and a clarion call for its naturalization. In much the spirit of Quine, they call for the marriage of metaphysics to science — in particular, to fundamental physics. Their tone is incendiary:

Their opening blast claims that contemporary analytic metaphysics 'contributes nothing to human knowledge': its practitioners are 'wasting their talents', and the whole enterprise, although 'engaged in by some extremely intelligent and morally serious people, fails to qualify as part of the enlightened pursuit of objective truth, and should be discontinued' (vii). (Dorr 2010, np)

Cian Dorr comments, “They set out on a 'mission of disciplinary rescue' in the spirit of Hume and the logical positivists, in which a fair proportion of philosophy as currently practiced... will be consigned to the flames” (2010, np). Ladyman and Ross argue that speculative metaphysics, to its detriment, relies for evidence on metaphysical intuitions. They argue that these intuitions are dubious because 1) they are culturally variable, 2) we have no evolutionary grounds for thinking they are truth-tracking, and 3) they have a poor historical track record (2007, 10 & 16). In contrast, naturalized metaphysics is essentially bound up with science. Naturalized metaphysics is, in their view, an attempt “to unify hypotheses and theories that are taken seriously by contemporary science” (2007, 1). Ladyman and Ross give explicit primacy to fundamental physics, such that its theses trump any conflicting theses from other sciences (2007, 44). Their central argument for the naturalization of metaphysics is that science has unmatched epistemic distinction, due to systematic error filters like rigorous peer review, which are repeatedly iterated and entrenched at the institutional level. Ladyman and Ross attempt to jar metaphysicians from their methodological complacency by 1) signalling the contaminating influence and epistemic bankruptcy of folk

intuition, and 2) defending a methodology respectful of the unmatched epistemic distinction of science, in which the task of metaphysics is to unify scientific theses.

Dorr is critical of the book in his *NDPR* review. First, he argues that Ladyman and Ross mischaracterize the role of intuition-talk in metaphysics. According to Dorr, they incorrectly assume that: “‘appeals to intuition’ are part of a distinctive *method* for doing metaphysics, a method we could contemplate giving up in its entirety” (2010, np). He argues that in fact, “‘intuition’ talk is playing no such distinctive role. Often, saying ‘Intuitively, P’ is no more than a device for committing oneself to P while signaling that one is not going to provide any further arguments for this claim” (2010, np). Talk of intuition is, he argues, often a stylistic choice, which introduces epistemic humility and avoids the flat-footedness of bald assertion. He explains:

one can make things a bit gentler and more human by occasionally inserting qualifiers like ‘it seems that’. It would be absurd to accuse someone who frequently gave in to this stylistic temptation of following a bankrupt methodology that presupposes the erroneous claim that things generally are as they seem... It may be bad style, but it is not bad methodology, or any methodology at all. (2010, np)

According to Dorr, intuition-talk is just talk; it does not signal a methodological epidemic. Further, Dorr comments that Ladyman and Ross “have missed what is best and most distinctive about the tradition they set themselves against: its gradual raising of the standards of clarity and explicitness in the statement of metaphysical claims” (2010, np). He concludes that if a desirable interaction is to take place between analytic metaphysics and modern science, then:

it will have to be pushed forward by philosophers with a foot in both camps, who combine a rigorous understanding of the space of interpretative possibilities opened up by the



physical theories with a metaphysician's patience for fine distinctions and quibbling objections. (2010, np)

That is, a fruitful interaction between them can only occur under the acknowledgment that the metaphysician has something to bring to the table.

Since *Every Thing Must Go* was published, there has been much second-order discussion of how to better articulate and motivate the call to naturalize metaphysics. Two important anthologies have been published on the subject: Chalmers, Manley, and Wasserman's *Metametaphysics* (2009) and Ladyman, Ross, and Kincaid's (2013) book, *Scientific Metaphysics*. From 2012-13, there was also an exchange on the topic of naturalizing metaphysics between Maclaurin and Dyke (2012, 2013) on the one hand, and McLeod and Parsons (2013) on the other.<sup>6</sup>

In their (2012) article, Maclaurin and Dyke try to develop and refine Ladyman and Ross' call for the naturalization of metaphysics. In response to Dorr, they acknowledge some benefits of non-naturalistic metaphysics, such as "achieving greater clarity" (2012, 301). They also explore a number of other potential sources of value (2012, 300-301). They respond to Dorr's claim that intuition talk is merely stylistic as follows:

We don't think that all the appeals to intuition that one finds in non-naturalistic metaphysics can be characterized in this way. Those that occur in Lewisian-style cost-benefit analyses, for example, do not fit this model. In such analyses two theories are compared with respect to a number of theoretical virtues, one of which is the preservation of common-sense intuitions. (2012, 296).

---

<sup>6</sup> Note that Maclaurin and Dyke's (2013) is listed in the works cited under Dyke to reflect first authorship.

Intuitions do, they argue, at least sometimes play an important evidential role in metaphysical theorizing. Non-naturalistic metaphysics, they say, is “any philosophical theory that makes some ontological (as opposed to conceptual) claim which, in principle, has no observable consequences” (2012, 292). It follows that naturalistic metaphysics makes ontological claims, *all of which* have observable consequences. Although they do not endorse Ladyman and Ross’ proscription of non-naturalistic metaphysics, they do encourage philosophers to take seriously their methodological questions and worries (2012, 304). Maclaurin and Dyke’s dubiousness of non-naturalistic metaphysics is based in part on the fact that, since non-naturalistic metaphysical theories lack observable content, we must select between them by “appealing to virtues that look more aesthetic than epistemic” (2012, 304). So our basis for selecting metaphysical theories has nothing much to do with “[the] aim of arriving at the objective truth about fundamental reality” (2012, 304). They also worry that non-naturalistic metaphysics doesn’t progress “beyond increasing the standards of clarity and precision in expressing its theories” (2012, 291). But their main criticism of non-naturalistic metaphysics is that “as an intellectual endeavour, it can have no practical benefit to anybody” (299). In particular, it does not benefit science (2012, 299).

McLeod and Parsons (2013) respond that Maclaurin and Dyke face difficulties similar to those of the positivists. According to McLeod and Parsons, Maclaurin and Dyke’s criterion for being naturalistic “closely resembles Ayer’s... weak verifiability criterion for ‘factualness’ of a proposition” (2013, 173). They argue that Maclaurin and Dyke’s notion of consequence is ambiguous and that “like Ayer’s criterion, Maclaurin and Dyke’s is trivial – either every theory counts as naturalistic, or no theory does” (2013, 173). That is because, on the one hand:

For every theory T, there's an auxiliary hypothesis of the form 'If T, then O', where O is an observation, which, when conjoined with T has the observation O as a consequence. So every theory counts as naturalistic according to Maclaurin and Dyke. (2013, 176)

All theories have observable consequences, on some auxiliary hypothesis. On the other hand, they argue, theories in science frequently fail to have observable consequences, in the sense of deductively implying them. This is shown by the fact that we frequently explain away failures of prediction, rather than rejecting the theory that produced them. For instance, when we discovered that measurements of the velocity of ether from different, moving locations, didn't capture motion relative to the ether, we didn't infer the falsity of the wave theory; rather, we supposed that motion relative to ether distorted the measurement apparatus to cancel out its effects (2012, 176-77). McLeod and Parsons conclude, "surely it's too high a standard to [require of] a philosophical theory that it must have empirical content in a way that even theories in the physical sciences fail to!" (2013, 177). Depending on how one cashes out the notion of *having an observable consequence*, the criterion for naturalism will either include too much or too little.

Dyke and Maclaurin (2013) respond by clarifying their notion of consequence.<sup>7</sup> They say that McLeod and Parsons point to a broader problem: "observations count for or against theoretical claims only on the assumption that the relevant auxiliary hypotheses are true" (2013, 180). We cannot demarcate science from non-science "unless we restrict the set of auxiliary hypotheses that scientists are allowed to employ" (2013, 180). Dyke and Maclaurin "interpret 'auxiliary hypothesis' in line with the practice of current science and philosophy of science" (2013, 180). Therefore, they accept only auxiliary hypotheses "best supported by current science" (2013, 180). They specify that when two different metaphysical theories, conjoined with the best auxiliary

---

<sup>7</sup> I have reversed their names to reflect first authorship.

hypotheses, “make exactly the same predictions about our observations... [t]his is what we mean when we call such theories non-naturalistic” (2013, 180).

Ladyman, Ross, Maclaurin, Dyke and their sympathizers may take heart, since naturalized metaphysics is on the rise. The emerging discipline of scientific metaphysics — which interprets and draws metaphysical conclusions out of science — is very rapidly establishing itself.<sup>8</sup> However, to borrow a fashionable term, the *metametaphysics* underlying this development remains disunified and inchoate. Its epistemological and methodological assumptions need to be spelled out, refined, and defended. How shall we distinguish naturalized from non-naturalized metaphysics? What makes a metaphysical theory naturalized or not? Why should we prefer one to the other? It’s here that I wish to enter the conversation.

### **III. *Physics-Venerating Verificationist Unificationism* Reconsidered**

#### **III.1 The Need for a Better Argument for Naturalization**

Like Maclaurin and Dyke (2012) and others, I contend that the task of articulating and defending the call to naturalize metaphysics didn’t end with Ladyman and Ross (2007). But Ladyman and Ross have been so seminal in the literature that I must say, explicitly, how and why I diverge from them. While I share Ladyman and Ross’ dubiousness of speculative metaphysics and their desire for metaphysics to be accountable to science, I’m compelled by neither their argument nor their program for the naturalization of metaphysics. I call their program of naturalized metaphysics *Physics-Venerating Verificationist Unificationism*, to capture their view

---

<sup>8</sup> See Calosi and Morganti (2016), Dorato and Morganti (2013), Esfeld (2004), French (2014), Kistler (2010), Loewer (2012), Maudlin (2007), Morganti (2008), Ney (2012), Ney and Albert (2013), Norton (2015), Pradeu and Guay (2015), Ruetsche (2011), and Waters (2017).

that naturalized metaphysics has one task: to unify theses that science verifies or declares verifiable, with explicit primacy given to physics. Let me first address their argument.

Ladyman and Ross argue that metaphysics should be naturalized because science has “epistemic supremacy” by virtue of its “repeated iteration of institutional error filters” (2007, 29). Science, they say, “just *is* our set of institutional error filters for the job of discovering the objective character of the world” and as such, science “will admit no epistemological rivals” (2007, 28). Their central example of these error filters is rigorous peer review. Science is “a community enterprise” and, unlike speculative forms of metaphysics, it is “not... supported by feats of individual reasoning” (2007, 28). Without rigorous peer review, Ladyman and Ross say that metaphysicians construct “a hermitically sealed world in which they can autonomously study their own special subject-matter” (2007, 14). So speculative metaphysics is not accountable to a community; science is.

I agree that the ability to identify, explain, and prevent errors is requisite for any discipline’s epistemic credibility. And it seems right that those errors should be filtered, actually or potentially, at the communal level. However, metaphysics *is* a community enterprise. As a philosophical enterprise, metaphysics is often a Socratic dialogue with others. Moreover, metaphysicians have peer review — their submissions for publication are double, sometimes triple-blind reviewed. Granted, metaphysical claims are not subject to the same *kinds* of error as scientific claims, so error filters can’t work for metaphysics like they work for science. For instance, one cannot falsify a grounding claim by pointing out that someone has miscounted or miscalculated. However, the claims of metaphysics are subject to other kinds of error, reflective of the sort of inquiry it is — namely philosophical inquiry. Its claims are subject to errors of reason: category mistakes, fallacies, explanatory deficiencies, counterexamples, and so forth. Further, the

arguments they constitute are subject to failures of validity and cogency. Finding such errors is one of the tasks of peer reviewers and a central activity of metaphysics more broadly. So it is false to say that the discipline of metaphysics lacks institutional error filters like peer review. Science is not unique in having peer review and other institutional error filters, so Ladyman and Ross simply haven't justified their claim that science admits no epistemic rivals.<sup>9</sup>

Furthermore, if, as Ladyman and Ross claim, “science is... demarcated from non-science solely by institutional norms [including] requirements for rigorous peer review” (2007, 28), and if metaphysics has the same or similar institutional norms, then Ladyman and Ross face a disastrous collapse. The distinction collapses between, on the one hand, the science to which they ascribe ‘epistemic supremacy’ and, on the other hand, the neo-scholastic metaphysics they ridicule and condemn. If the distinction between science and neo-scholastic metaphysics collapses, Ladyman and Ross’ entire project collapses with it.

Ladyman and Ross might respond that peer review in metaphysics fails to be sufficiently *rigorous*. But then, for their argument for naturalization to be complete, they must give criteria for the appropriate kind of rigour, such that science meets the criteria and neo-scholastic metaphysics does not. Since they give no such criteria, their argument is incomplete. At worst, Ladyman and Ross’ defence of naturalized metaphysics entirely collapses; at best it is incomplete.

### **III.2 The Need for a Better Program of Naturalization**

I also take issue with several aspects of Ladyman and Ross’ program of naturalization. I explain why in the following sections.

---

<sup>9</sup> That’s not to say that their conclusion is false — just that their argument is unsound.

### III.2.1 Verificationism

Ladyman and Ross' program of naturalization builds in verificationism. Their verificationism is 'non-positivist' (2007, 29) since it is not an account of meaning, but rather, of being a contributor to objective inquiry (2007, 30). Their verificationist criterion is as follows:

No hypothesis that the approximately consensual current scientific picture declares to be beyond our capacity to investigate should be taken seriously... [and] any [serious] metaphysical hypothesis... should have some identifiable bearing on the relationship between at least two relatively *specific* hypotheses that are either regarded as confirmed by institutionally *bona fide* current science or are regarded as motivated and in principle confirmable by such science. (2007, 29)

They explain that to have 'bearing on' is 'to motivate' (2007, 30) and that a proxy for specificity is fundability (2007, 34 & 38). For Ladyman and Ross, if something is to contribute to objective inquiry, it must be verifiable in that sense.

*Prima facie*, Ladyman and Ross' verificationism undermines the very scientism underlying their call to naturalize metaphysics (see 2007, Chapter 1). It may undermine certain aspects of science, since science declares many of its own theses to be 'beyond our capacity to investigate'. Take, for example, the theses of string theory. Richard Dawid explains: the theory potentially allows for  $10^{500}$  up to  $10^{1000}$  ground states — or lowest-energy states — and “a huge number of local ground states suggests that a correspondingly huge number of low energy patterns of parameter values are physically possible” (2013, 86). Because it is impossible to scan so many ground states, we cannot confirm that “the universe we live in corresponds to a valid [set of] string theory ground state[s]”, and “[i]n this light, it is difficult to assess how predictive string theory can be in principle” (2013, 86). So Ladyman and Ross' verificationist criterion for being a contributor

to objective inquiry is too restrictive for their purposes, since even science sometimes fails to meet it.

It is difficult to say how Ladyman and Ross might respond, since their verificationist criterion appeals to ‘the approximately consensual current scientific picture’ — and here, the scientific community is far from showing consensus. In fact, it is a matter of much current controversy among physicists whether string theory has testable empirical consequences and hence whether it counts as genuine science (Ellis and Silk 2014, Smolin 2006, Woit 2006, Wolchover 2015). Ladyman and Ross don’t say what to do when we lack consensus. They acknowledge some of the criticisms of string theory (2007, 168-70), but do not weigh in on the matter. They could take a hard line and deny that string theory follows the institutional norms that demarcate science from non-science and so fails to be science. Or they could claim that while string theory counts as science, it can’t contribute to objective inquiry.<sup>10</sup> But the better option is for them to remain agnostic on these matters and hope that some consensus emerges in the scientific community about string theory’s empirical and scientific credentials. If Ladyman and Ross take this line, then their verificationism *need not* undermine their scientism. If consensus emerges in the scientific community about the empirical and scientific status of string theory (and other controversial bits of science), they can simply adopt the consensus view; if no consensus emerges, they can remain agnostic.

A second, more decisive problem is that the failure of Ladyman and Ross’ demarcation criterion renders untenable their verificationist criterion for being a contributor to objective inquiry. That’s because the criterion hinges on a demarcation between “*bona fide* current science”

---

<sup>10</sup> This line may sound odd, but it is available to them, since they don’t explicitly build their verificationist criterion into their demarcation criterion, and since they claim that some sciences are not objective inquiries (2007, 36).



and other forms of inquiry (2007, 29). I have already suggested that their demarcation fails. If it does, then their verificationism fails, too. And perhaps we shouldn't be surprised. The verificationists famously failed, despite great effort, to satisfactorily formulate their verificationist criterion of meaning (see Lewis 1988; McLeod and Parsons 2013; Psillos 1999). As Lewis explains: "The collapse of Ayer's criterion, and then the sorry history of unintuitive and ineffective patches, have done a lot to discredit the very idea of delineating a class of statements as empirical" (Lewis 1988, 4). It's not clear why a verificationist criterion for being a contributor to objective inquiry should fare any better. If those sounding the call to naturalize metaphysics take up the verificationist mantle, they must contend with the historical failures that have been taken to decisively show its bankruptcy.

### **III.2.2 Modes of Engagement**

The second feature of Ladyman and Ross' program of naturalization that I wish to take issue with is the narrow role they assign to naturalized metaphysics. Ladyman and Ross assign only a single task to naturalized metaphysics: the unification of scientific theses (2007, 1). They impose on naturalized metaphysics the following principle, which they call the Principle of Naturalistic Closure:

Any new metaphysical claim that is to be taken seriously at time  $t$  should be motivated by, and only by, the service it would perform, if true, in showing how two or more specific scientific hypotheses, at least one of which is drawn from fundamental physics, jointly explain more than the sum of what is explained [by them individually]. (2007, 37)

But a naturalized metaphysics need not engage so narrowly with science. In fact, metaphysics can engage with science in many ways, some of which I will spell out in Chapter 3. For now, the

important point is that a program for naturalizing metaphysics shouldn't be *unnecessarily* restrictive regarding the role it assigns to metaphysics in relation to science. The more metaphysics *can* do for us, the more we should ask it to do.

### III.2.3 Unificationism

Ladyman and Ross hinge the possibility of naturalized metaphysics on the possibility of unification in the sciences. They claim that “one of the important things we want from science is a relatively unified picture of the world” (2007, 27). In their view, naturalized metaphysics helps assemble that picture: “the *raison d'être* of a useful metaphysics is to show how the separately developed and justified pieces of science (at a given time) can be fitted together to impose a unified world-view” (2007, 45). Ladyman and Ross explicitly acknowledge that if unification were not possible, then there would be no role for naturalized metaphysics: “if the world were fundamentally disunified, then discovery of this would be tantamount to discovering that there is no metaphysical work to be done” (2007, 5-6). For clarity, what they should say is not *if the world* were fundamentally disunified, but if *science* were.

But it is increasingly common to think that science *is* fundamentally disunified. Philosophers of science are moving away from the view that unification is inevitable or possible, toward various scientific pluralisms (Chang 2012; Kellert, Longino, and Waters 2006; Kitcher 2001; Mitchell 2002, 2003, and 2009; Ruphy 2017). Science is messy and unification elusive (Dupré 1993, Neumann 1978, Rosenberg 1994, Ruetsche 2011). Why make the possibility of naturalized metaphysics contingent on such a dubious prospect as scientific unification? Why make the limit of unification the limit of naturalized metaphysics? A program for the naturalization of metaphysics should reflect the disunified reality of scientific theory and practice.

### III.2.4 Physics Primacy

Lastly, in their program for the naturalization of metaphysics, Ladyman and Ross assign explicit primacy to physics (see 2007, Chapter 1, §1.4). Their program of naturalization “accords physics a special status”, such that “for a metaphysical claim to be taken seriously it must relate to at least one specific scientific hypothesis of fundamental physics” (2007, 39). And, in their view, physics is the trump card. That is, they include in their program a Primacy of Physics Constraint, according to which “[s]pecial science hypotheses that conflict with fundamental physics, or such consensus as there is in fundamental physics, should be rejected for that reason alone. Fundamental physical hypotheses are not symmetrically hostage to the conclusions of the special sciences” (2007, 44). Ladyman and Ross don’t motivate physics primacy by appeal to reductionism. They deny that the world “comes in ‘levels’” (2007, 54), so they must deny that one level reduces to another. Instead, they argue that over the history of science, physical theory:

has been unified and extended. Consolidations and unifications... are part of the reason for supposing that there is a coherent body of fundamental physical theory of sufficient scope and power that it is the only candidate for the ‘most basic and comprehensive of the sciences’. (2007, 44)

Physics, they think, has primacy not because everything reduces to the physical, but because the discipline of physics has the greatest scope among the sciences. But what exactly does physics tell us? If we look for an approximately consensual theory, we do not find one — we find, rather, a bewildering array of discontinuous theories and theoretical interpretations (Ruetsche 2011). Once we boil physics down to its bare consensus bones, would it still have sufficiently great scope to justify its primacy over other sciences? And if the pluralists are right and the sciences are likely to

remain disunified, then shouldn't naturalized metaphysics proceed through piecemeal engagement with various sciences, theorizing about various levels of organization? And if science doesn't unify, we might not expect metaphysics to, either. In short, we need not venerate physics over the other sciences. That is, we need not be so *singularly* concerned with physics.

### **III.3 Desiderata for the Call to Naturalize Metaphysics**

By way of summary, let me syphon from my criticisms of Ladyman and Ross some desiderata for arguments and programs for naturalizing metaphysics. First, the arguments and programs must delineate clearly between science and non-naturalistic metaphysics if they are to assign distinction to the one and criticize or denounce the other. They must also avoid the failures of historical campaigns against metaphysics, such as logical positivism and its accompanying verificationism. Further, programs for the naturalization of metaphysics shouldn't assign a needlessly narrow role to metaphysics in relation to science. They should also reflect the complicated realities of scientific theory and practice, by giving a role to metaphysics even if pluralism is correct and unification isn't possible. Finally, being friendly to pluralism, would-be naturalizers shouldn't insist on the primacy of physics, since engagement with various forms of science may be fruitful and since pluralism in science may motivate pluralism in metaphysics. What I have called Ladyman and Ross' *Physics-Venerating Verificationist Unificationism* doesn't respect these desiderata and is therefore dissatisfactory. Work is needed to better motivate and refine the call to naturalize metaphysics.

Let me briefly foreshadow the ways that my project here will respect the desiderata I have outlined. First, since my demarcation of science from non-science was ostensive, it does not threaten to collapse. Second, I will disavow all forms of verificationism and thereby avoid the

problems that historically plague its articulation. My program of naturalization will not build in verificationism of any kind, since it will not require that every metaphysical thesis be empirically tractable. It will make empirical content a necessary condition neither for a claim's meaning, nor for its interest and importance. Rather than demanding that each metaphysical thesis be empirically or scientifically verifiable, I will suggest only that metaphysical theories as a whole should be motivated and supported by scientific evidence. Further, I will not narrowly restrict the role of metaphysics in relation to science, but instead outline a variety of ways that scientifically responsible metaphysics may engage science. Further still, on the conception of scientifically responsible metaphysics that I will articulate, the impossibility of unification in science and the truth of pluralism would not put metaphysics out of work. Unification is not, in my view, the sole task of scientifically responsible metaphysics. Finally, I will leave the door open for metaphysical engagement with various scientific theories, interpretations, and practices, without assigning primacy to physics.

#### **IV. Armchair Metametaphysics?: Methodological Considerations**

This is a project in second-order metaphysics, or what is frequently called *metametaphysics*. It poses epistemological and methodological questions about metaphysics, as opposed to posing first-order metaphysical questions about the world. But here one might ask: aren't there *further* epistemological and methodological questions to be asked about that second-order inquiry? Aren't there *third-order* questions, regarding which epistemological principles and methodological practices we should accept for projects like this one?<sup>11</sup> If so, then when I consider

---

<sup>11</sup>Note the threat of infinite regress.

the proper relationship of first-order metaphysics to current science, I should also consider the proper relationship of second-order metaphysics to it.

Now, one might think that what goes for the one goes for the other — that first-order and second-order metaphysics stand or fall together. One might think that whatever arguments I give against non-naturalized metaphysics will also undermine non-naturalized *metametaphysics*. Likewise, whatever arguments I give for the preferability of naturalized metaphysics will also demonstrate the preferability of naturalized *metametaphysics*. And, the thought goes, if those arguments are not themselves based on science, they will ultimately be self-undermining. In other words, if one reaches the conclusion ‘naturalize’ using methods that take no account of science, then one cuts off the branch on which one sits. If so, then the call to naturalize must itself be naturalistic.

But first and second-order metaphysics need not stand or fall together. The two differ importantly, since their domains differ. First-order metaphysics is about *the world*; second-order metaphysics is about *metaphysics* — its possibility, its subject-matter, and its methods. Further, while the subject-matter of first-order metaphysics is *descriptive*, the second-order metaphysics that I do here is *normative*. Now, it is possible — maybe even to be expected — that we should appeal to substantially different types of evidence for inquiries with substantially different subject-matters. The kinds of evidence fruitful in one domain may not conduce to knowledge in other, quite different domains.

While science is a good source of evidence about the world, what comparable source of evidence do we have about the proper methods of metaphysics? One might think that while the proper methods of metaphysics don’t fall within the *domain* of science, science can still inform us about them — it can lead by example. Since the sciences constitute our most successful ways of

knowing, all inquirers, including first-order and second-order metaphysicians, should follow whatever epistemological and methodological principles tacitly underlie scientific practice.<sup>12</sup> Call this the *Maximal Scope* view, since it holds that we should apply the principles underlying scientific practice in all cases. The Maximal Scope view is really the conjunction of epistemological and methodological naturalism. If the Maximal Scope view is right, then this dissertation should identify the principles underlying various scientific practices and explain how they apply in first-order metaphysical practice.

But the Maximal Scope view is false. Again, it's not clear that we should we lump all forms of knowledge together and suppose that the principles conducive to knowledge in one domain will conduce to knowledge in substantially different ones. But even if we suppose that our epistemological and methodological principles for these domains will be uniform, *science simply does not hand us a normative epistemology*. It may, debatably, hand us an epistemology, but we can't get the normativity without some *a priori* assumptions. Let me explain.

Suppose we describe the methodological and epistemological principles underlying some scientific practices. Doing so would give us some descriptive facts. But how do we get from those descriptive facts to the *normative* claim that we should adopt the relevant principles in other domains? Here we face the is-ought gap. While Prior (1960) showed that the gap doesn't hold when we exploit the way that truth preservation works, e.g. by introducing normative claims as disjuncts, descriptive statements don't imply normative ones in usual cases, i.e. cases where the normative claim is simple (Gibbard 2012) or relevant (Singer 2015).

---

<sup>12</sup> Note that this stops the threat of regress. There is no need for proliferating meta-inquiries, since there is just one set of epistemological and methodological principles, applicable in all cases.

That's not to say that you *can't* get an ought from an is. You can — but you need a bridge principle to do it. You need a principle stating 'if  $x$  is the case, then  $y$  ought to be the case'. Let's look at an example. The Maximal Scoper could argue something like the following, where 'BP' signals the bridge principle:

P1) Epistemological and methodological principles  $\{\alpha, \beta \dots n_1\}$  underlie scientific practices  $\{x, y \dots n_2\}$ .

P2) Principles  $\{\alpha, \beta \dots n_1\}$  conduce to the great empirical success of practices  $\{x, y \dots n_2\}$ .

BP) If some epistemological and methodological principles conduce to great empirical success for some practices, then we should adopt them in other domains.

---

C) We should adopt principles  $\{\alpha, \beta \dots n_1\}$  in other domains.

The bridge principle gets us from the descriptive premises to the normative conclusion.

Now, what justifies the bridge principle? It is justified by a further, implicit premise:

IP) Great empirical success is an epistemic good.

Of course, it is part of the job of a normative epistemology to identify epistemic goods. But what justifies the implicit premise? It identifies some criterion, or end, or in Quine's words, "terminal parameter" that we wish to promote (qtd. in Kornblith 2014, 72). Naturalistic epistemologies that aim to preserve normativity must begin with such terminal parameters in order to establish the epistemic credentials of science. For instance, they must claim that science is *empirically successful, truth-conducive, subject to Bayesian confirmation, efficacious* with respect to certain pragmatic purposes, and so on. For Quine, the terminal parameter is "efficacy for an ulterior end, truth or prediction" (qtd. in Kornblith 2014, 72). In his view, a naturalistic epistemology can preserve normativity by expressing terminal parameters, then descriptively studying how those parameters are met; naturalistic epistemology is "the technology of truth-seeking" (qtd. in



Kornblith 2014, 71). But as Kornblith points out, there isn't consensus, even among naturalists, about the terminal parameter (2014, 72). Even if there were, there would still be the additional question of what the consensus *should be*. Whatever our terminal parameter, the question remains: why *that* parameter?

The choice of terminal parameter can never be made from epistemologically neutral ground and must ultimately be *a priori*. If we take science to be a good guide to epistemology, we've judged it to be epistemically meritorious, but to do so, we must have already made some epistemological assumptions — reliabilist, pragmatic, Bayesian, or whatever. That's not to say we begin with a fully worked-out epistemology — we begin only with some intuitive assumptions about what is epistemically good. Since we can only get normativity into the foundations of epistemology by relying on those intuitions, normative epistemology must always bottom out in the *a priori*. If so, then scientific evidence does not bear on the foundations of epistemology, and the Maximal Scope view must be false. If epistemology includes a normative dimension, then there can't be a fully naturalized epistemology.<sup>13</sup> And finally, normative second-order metaphysics cannot simply appeal to science.

One might object: but once we've identified some terminal parameter, the question of whether metaphysics meets it *can* be settled by scientific inquiry. Suppose that we settle on true belief as the ultimate epistemic goal. Having settled that, then why shouldn't we simply look to

---

<sup>13</sup> Compare Putnam (1982), who argues that the attempt of the naturalized epistemologist to eliminate the normative is futile, because a non-reductive and fundamentally normative notion of truth is indispensable:

“Let us recognize that one of our fundamental self conceptualizations... is that we are thinkers, and that as thinkers we are committed to there being some kind of truth, some kind of correctness which is substantial and not merely ‘disquotational’. That means that there is no eliminating the normative” (1982, 21).

the history of metaphysics, delineate the various methods it contains, and check which of them have been most conducive to true belief. This would arguably be a kind of scientific study. On this view, the history of metaphysics is a history of methodological trial and error, and we must simply determine what have been the successes and failures.

If such a project can be undertaken fruitfully, then I'm in favour of it. However, I doubt that it can. For many metaphysical theories and claims, it is difficult if not impossible to assess their truth. One reason is that, as we have seen, we don't get the same kind of error-correction in metaphysics that we do in science. Our metaphysical theories tend not to make novel predictions, so the world typically can't frustrate them – this is partly why there are Platonists two-thousand years after Plato. The other barrier to assessing the truth of metaphysical claims is that many metaphysical claims are, as the positivists pointed out, not verifiable. Sometimes it is because they are couched in language that prevents truth-evaluation. As David Chalmers argues, because many metaphysical terms are ill-defined, vague, or ineliminably intuitive, metaphysical claims frequently “lack determinate truth-conditional content, and typically lack determinate truth-values” (2009, 127). For instance, Chalmers thinks we do not have a grip on what it would mean for a table or a number to *really exist* (2009, 103). The point is that metaphysical discourse sometimes relies on obscure philosophical terms for which there is no consensus on meaning or correct usage; for which there is no precise analysis; and which, as a result, preclude metaphysical claims from having determinate truth-conditions. Sometimes the problem is not with the language, but with the kind of claim being made. Some metaphysical claims – claims about the forms or about God – are about things that we, in principle, can have no direct knowledge of (and, in the case of the forms, no causal contact with). Given our limited epistemic circumstances (barring extraordinary things like divine revelation or communion with the forms), it doesn't seem we can

verify such claims. That's not to say we cannot have *evidence* for them – just that we cannot establish their truth. In sum: many metaphysical claims are not truth-evaluable due to the language in which they are couched or due to their subject-matter. While I don't *rule out* a project that aims to catalog the methodological successes and failures of metaphysics, I think many metaphysical claims have features that would frustrate that project. So I won't undertake that project here.

In sum: this section considered the claim that if first-order metaphysics should be naturalistic, then so too must second-order metaphysics, since the two stand or fall together. I argued that they don't stand or fall together, since their subject-matters differ substantially. So the epistemological and methodological principles that go for the one need not go for the other. I considered and rejected the Maximal Scope view, according to which all inquirers should adopt the epistemological and methodological principles underlying scientific practice. First, I suggested that substantially different subject-matters may call for substantially different methods. Second, I argued that the Maximal Scope view can only get its normative dimension on an *a priori* basis, which entails the falsity of the view. I considered and expressed skepticism about the prospects of the view that we could just do a scientific study of which metaphysical methods have been most truth-conducive. The upshot for this dissertation is that I may proceed without explicit appeal to science.

## **V. Plan for the Dissertation**

Let me now outline the plan for the rest of the dissertation. In Chapter 2, I will articulate and defend the epistemological principles on which my evaluation of, and prescriptions for, metaphysics rest. The aim of the chapter will be to argue that for epistemic purposes, as opposed to heuristic or pragmatic purposes, theories should be robustly constrained and adequately

warranted. I will situate the chapter in relation to the epistemic normativity literature, articulate two sets of epistemic principles, and argue for those principles.

Equipped with the epistemic norms established in Chapter 2, in Chapter 3, I will show how they apply to contemporary metaphysics. Contrary to the provocative sentence with which I opened this chapter, I won't quite argue that analytic metaphysics is bunk. I will argue, rather that some of it is epistemically inadequate, and that what I call *scientifically responsible metaphysics* is preferable for the purposes of producing justified theories about the world. I will criticize what I call *free range metaphysics* — metaphysics, the content of which is not constrained by science, but rather, by logical, aesthetic, and psychological demands, such as the demands for consistency, simplicity, intuitive plausibility, and explanatory power. I will show how the constraints on such theories individually and jointly fail to be sufficiently robust and to secure adequate epistemic warrant. I will then argue that a scientific constraint conduces to robustness and epistemic warrant. That is because the disconfirmation of scientific theses creates robust theoretical constraint, and their confirmation secures epistemic warrant.

Chapter 4 will consider potential problem cases — cases of metaphysical topics not *obviously* apt to be made scientifically responsible. However, I will resolve the problem cases by showing how the topics can be made scientifically responsible. In doing so, I will demonstrate what scientific responsibility looks like on the ground. Part I will consider modal metaphysics, which is problematic insofar as *prima facie*, actuality does not acquaint us with possibility or necessity. Nevertheless, the chapter will articulate a scientifically responsible methodology for modal metaphysics that takes current science as an evidence-base for the justification of modal claims and as a model of good modal reasoning. Scientifically responsible modal metaphysics does not rely on dubious metaphysical intuitions and betters purely rationalist modal metaphysics to the

extent that it allows for greater resolution of disagreement, accounts for modal error, and provides modal metaphysics with much needed discipline and control. Part II will deal with the grounding literature, which is *prima facie* anti-naturalistic, insofar as it rejects the Quinean view of metaphysics in favour of an Aristotelian one and posits ontological dependence relations seemingly alien to science. I will perform a piecemeal analysis of various influential works in the grounding literature, starting with explicit second-order commitments and proceeding to first-order practices. In my examination of the grounding literature, I will uncover a heterogeneous array of methods. I will synthesize a list of fruitful uses of science for grounding theorists, including among other things: to help in the identification of putative grounding relata, to show correlations among them, to demonstrate their non-identity, to provide a stock of explanatory patterns, to identify candidate essential properties, and to motivate agnosticism about particular grounding theses where scientific support is lacking. I will conclude that the prospects for making both modal metaphysics and theories of ground scientifically responsible are bright.

Overall, the project has considerable importance with respect to the discipline of metaphysics, since it will show that large swaths of metaphysical inquiry are epistemically inadequate. It will show an impetus for reform and spell out the proper mode of reform. It will provide would-be naturalists with a guide for fruitful engagement with science, by filling out an account of scientifically responsible metaphysics that is, in comparison with other calls to naturalize, more measured and reflective of the complicated reality of scientific practice.

## CHAPTER 2

### Epistemic Principles

In this chapter, I will articulate and motivate the epistemic principles on which my evaluation of, and prescriptions for, contemporary metaphysics will rely — principles revolving around theoretical constraint and epistemic warrant. The aim of the chapter is to argue that for epistemic purposes — as opposed to heuristic or pragmatic purposes — theories should be robustly constrained and adequately warranted. In section I, I will introduce the aims of this chapter with reference to the epistemic normativity literature. I will suggest that, just as there is an ethics of belief, there is something like an ethics of theory-building. One way of characterizing the central aim of this chapter, then, is to supply some of the content of that ethics of theory-building. In section II, I will articulate the two sets of principles that I wish to motivate. In section III, I will defend the principles. I conclude that I may proceed with right, in Chapter 3, to evaluate theories based on whether they are robustly constrained or sufficiently warranted.

#### I. Epistemic Normativity

The ethics of belief literature is committed to the view that we morally ought or ought not to believe certain things under certain conditions. That is, it commits to there being ethical norms regarding the adoption of beliefs. Suitably adapted meta-ethical theories have been used to account for those norms, including deontology, consequentialism, and virtue ethics.<sup>14</sup> Søren Klausen notes

---

<sup>14</sup> On epistemic deontology, see BonJour (1985), Chisholm (1989), Feldman (1988), Ginet (1975), Moser (1989), Nagel (1986), Stapleford (2015). On epistemic consequentialism, see Ahlstrom-Vij and Dunn (2014), (forthcoming); Briesen (2016); Goldman (2015); Klausen (2009); Petersen (2013) and for criticism see Berker (2013) and Littlejohn (2015) and (forthcoming). For virtue-theoretic accounts of epistemic normativity, see Axtell (1997), (2000);

that deontology has been the “overwhelmingly dominant” view in epistemology (2009, 162). At any rate, according to these views, given some favourable epistemic circumstances, we morally ought to believe  $p$  because it is an epistemic duty, or because it maximizes some epistemic good, or because it is epistemically virtuous. There is much room for disagreement within each theoretical framework — for instance, we might disagree about the epistemic good (or goods) that our actions should preserve or promote. Still, the common thread that unites the views is their commitment to there being ethical norms governing belief.

We might adopt a parallel view with respect to theories. Theories and belief structures share a good deal in common, after all. They consist of a number of interdependent epistemic commitments that are, ideally, roughly consistent. We hope that our beliefs enable us to understand and navigate the world; likewise we hope that our theories help us to explain, predict, and successfully intervene on bits of the world. Of course, the analogy is not perfect. While a person’s set of beliefs might be called a theory, it is certainly a different *kind* of theory than a philosophical or scientific theory. A set of beliefs is not about some particular subject-matter. Rather, I have beliefs about all sorts of things: what I had for breakfast, which politician deserves my vote, which is the best flavour of ice-cream, where I am presently, and so forth. Theories, on the other hand, tend to be about some particular subject-matter: the physical underpinnings of the world, the motions of the planets, and so forth. Theories also aim for a level of abstraction or generality that is explanatory with respect to some particular body of evidence; to that end, they frequently consist of general propositions. Beliefs, on the other hand, don’t typically aim toward generality and are often expressible as singular propositions. Further, however socially shaped beliefs may be, they

---

Axtell and Carter (2008); Greco (2010); Henning and Schweikard (2013); Montmarquet (1993); Sosa (2007); Thomas (2008); Zagzebski (1996); Zagzebski and Fairweather (2000).

are personal; the sorts of theories we are interested in here are products of disciplinary activity. Notwithstanding these differences, sets of belief and theories are similar sorts of thing.

If the analogy holds, then just as there is an ethics of belief, there is something like an ethics of theory-building. I say ‘something like’ because I do not wish to commit to the relevant sort of normativity being *moral*. There might be a case for thinking the relevant norms have a specifically moral dimension, parallel to the case for there being specifically moral norms governing belief — a case that stresses how theories impact real-world decisions, policies, and so forth. But the relevant oughts might well be epistemic oughts, rational oughts, or some other kind. I will remain neutral about this. For my purposes, it is not important *which* kind of normativity governs theory-building — the crucial point is just that *some* kind does.

Moreover, I assume that whatever kind of *ought* applies to theory-building, that ought implies *can*. In other words, if it’s true that we should follow a certain method in our theorizing, then we must be able to follow that method. In other words, the sorts of norms I am interested in are not norms for agents with unlimited cognitive, epistemic, and pragmatic resources, but rather norms for human inquirers with limited resources.<sup>15</sup> None of the norms I will prescribe demand anything of inquirers that is beyond their capacities.<sup>16</sup>

The possibility of an ethics of belief may hang on the truth of *doxastic voluntarism* — that is, on our ability to choose what we do and don’t believe.<sup>17</sup> If it’s true that I should believe *p* given

---

<sup>15</sup> Compare Cherniak's (1986) notion of minimal (as opposed to ideal) rationality, which indexes the conditions of rationality to our limited human capacities.

<sup>16</sup> Not only should the norms not demand things of us that are beyond our capacities, but they should be appropriately suited to our capacities as well. I acknowledge that different norms might be appropriate to creatures with different capacities. I thank Michael Devitt for raising this point.

<sup>17</sup> See Adler (2002), Alston (1985) and (1988), Audi (2001), Feldman (2001), Ginet (2001), Jäger (2004), Owens (2000), Ryan (2003), Shaffer (2013), Steup (2000) and (2008), Weatherson (2008); in response see Chrisman (2008), Hieronymi (2008), and McHugh (2012).



my evidence, it must be the case that I can elect to believe  $p$  — that is, if I *ought* to believe it, I must be *able* to believe it. However, the possibility of an ‘ethics’ of theory-building doesn’t likewise hang on the truth of doxastic voluntarism. That is because doxastic voluntarism is a view about the ability of individual people to choose which things they believe. But the ‘ethics’ of theory-building does not concern *personal beliefs*; it concerns *theoretical claims*. Moreover, the norms it posits do not concern the act of *believing*, but of *accepting* certain theoretical claims (in a sense I will define below — in short, adopting a policy of making the claims). Lastly, the ‘ethics’ of theory-building is not directly about the normative requirements on *individual believers*, but on *disciplinary communities*. So the possibility of an ‘ethics’ of theory-building does not hang on the truth of doxastic voluntarism as it does in the case of the ethics of belief, since the one has no bearing on the other.

Whichever approach we take to our ‘ethics’ of theory-building — deontic, consequentialist, virtue-theoretic — a more fundamental issue remains. As Klausen explains, “[d]iscussions of epistemic normativity are usually about finding the *basic* normative factor – the *ultimate* goal or value – in the epistemic domain” (my emphases, 2009, 161). That is, we must decide what we take to be the ultimate epistemic good or goods — the good that grounds epistemic duty, that constitutes epistemic utility, or toward which the epistemically virtuous agent directs her activities and attitudes. Here I understand *final epistemic value* to mean *intrinsic* epistemic value. However, there are alternate ways of cashing out this notion of finality. For instance, Grimm (2009) gives an account in terms of socially shared value. At any rate, the received view, as Klausen puts it, is that

the ultimate epistemic good is true belief (2009, 161).<sup>18</sup> The view that true belief is the ultimate epistemic goal is known as *veritism*.<sup>19</sup>

Veritism is *prima facie* plausible. What other epistemic goods *are there* that could compete with true belief as a candidate for the ultimate epistemic good? It turns out there are many. Internalists argue, against veritism, for the greater value of *justified* true belief over true belief (DePaul 2001, Feldman 2002, Smithies 2012). Others argue that, rather than justified true belief, *knowledge* is the final epistemic aim.<sup>20,21</sup> Klausen lists several more candidates for ultimate epistemic value: rational acceptability, agreement, fitness or pragmatic success, and the avoidance of falsehood (2009, 161). Note that several authors (Klausen 2009, Whiting 2013b, and Zagzebski 2003) argue that attaining true beliefs and avoiding false ones are *distinct* epistemic values. Linda Zagzebski (2003) argues further that the two aims can come into conflict. At any rate, Jeffrey Dunn lists several more candidates for final epistemic good: “the informativeness or interestingness of the propositions believed... whether the propositions believed are mutually explanatory or coherent... wisdom (Whitcomb 2007), understanding (Kvanvig 2003), or a love of truth (Zagzebski 2003)” (2016, np).<sup>22</sup> Further, Clayton Littlejohn (2015, forthcoming) argues for the

---

<sup>18</sup> Klausen actually says that the received view is that *truth* is the ultimate epistemic good. I take this to be shorthand for true belief.

<sup>19</sup> See Ahlstrom-Vij (2013), Coates (2009), David (2013), Goldman (1999) and (2002), Kornblith (1993), Lynch (2004), Steglich-Petersen (2013), Sylvan (forthcoming), and Whiting (2012), (2013a), and (2013c).

<sup>20</sup> See Adler (2002), Bird (2007), Engel (2007), Huemer (2007), Littlejohn (2013), McHugh (2011), Peacocke (1999), Sutton (2007), Williamson (2000); in response, see Baehr (2009) and Frise (2017).

<sup>21</sup> I distinguish knowledge from justified true belief because, beyond the fact that the Gettier cases show the joint insufficiency of the JTB conditions, several authors in the literature think justification is not necessary for knowledge (see Goldman 1967, Armstrong 1973, and Nozick 1981).

<sup>22</sup> Note that for Zagzebski (2003), love of truth does not have *final* epistemic value in the sense of its value being independent of any other sources of value, since its value is partly dependent

falsity of veritism on the grounds that the motivational role of belief can be a fundamental epistemic good. Additionally, Stephen Stich (1993) argues that the criterion by which we judge epistemic practices should be their aptness to promote what we value. Given this array of competing views, it is not clear that we should expect there to be just one final epistemic value — there might be two or more (Wright 2014) or a kind of epistemic value pluralism (Brogaard 2008, DePaul 2001, Weiner 2014; in response, see Ahlstrom-Vij 2013). Regardless of whether we go monistic or pluralistic regarding final epistemic value, it is clear that we have several rival candidates for ultimate epistemic good.

So we face a choice at the very foundations of epistemology. The choice regards which epistemic good or goods we take to be fundamental. I raise this because one *might* think that the question of the fundamental epistemic good must be settled before I can proceed with the present project. After all, I aim to motivate some epistemic norms and one's account of epistemic normativity will ultimately hinge on what one takes to be the ultimate epistemic good. In particular, how to flesh out and decide between deontic, consequentialist, or virtue-theoretic accounts of epistemic normativity depends on what we take the final epistemic value(s) to be. Moreover, assignments of epistemic distinction to forms of inquiry — including science — sometimes assume an answer to this foundational question. For instance, recall from Chapter 1 that Quine assigns epistemic distinction to science because it successfully promotes his “terminal parameter” — namely, “truth or prediction” (qtd. in Kornblith 2014, 72). As I argued in Chapter 1, I take this to undermine the possibility of a fully naturalized epistemology. At any rate, one might think that in order to get epistemic normativity off the ground in the first place and in order to assign epistemic

---

on the value of true belief. But in her view, love of truth confers “additional value on the acts it motivates” (2003, 147). Compare Hazlett (2013), who argues that true belief has no value independently of love of truth.

distinction to science, I must first settle the foundational epistemological question – or at least assume some answer to it.

However, I will set aside the epistemological foundations for my purposes here.<sup>23</sup> I can do so because I will remain neutral about the nature of epistemic normativity as such and because, while defenders of scientism and more moderate pro-science attitudes sometimes assume a final epistemic value, they *need not* do so. Rather, they can simply claim that science promotes some epistemic good, without claiming it to be final. For instance, I take Ladyman and Ross (2007) not to make any claims about the final epistemic good. What I aim to do, quite apart from settling the foundational epistemological question, is to establish some of the contents of the ‘ethics’ of theory-building. That is, I aim to articulate and motivate some normative epistemic principles that govern the construction of theories. To motivate my epistemic principles, I do not need to take a stance on the foundational question. I just need to show that the principles promote some epistemic goods that we frequently do care about, regardless of whether or not those goods are fundamental. As such, my conclusions in this chapter will be conditional on our recognizing certain epistemic goods. But I will show that the same principles hold for a number of different epistemic goods. In particular, if we take greater statistical likeliness, agreement, avoidance of falsity, methodological expediency and systematicity, truth-conduciveness, maximization of relevance, or rationality to be epistemic goods, then we ought to adopt at least some of the epistemic principles I will articulate in section II below. My defence of the epistemic principles in section III will consist in my showing how they promote each of those goods.

---

<sup>23</sup> I will at points make *other* sorts of foundational claims, e.g. about the aims of metaphysics. So I am not setting aside all foundational issues — just these questions about ultimate epistemic value.

## II. Epistemic Principles

### II.1 Theoretical Constraint

In my view, one of the features of theories that we should care about is the extent to which their content is *constrained*. A theoretical constraint is a limit on which theoretical contents we take to be admissible into a given theory. When we accept, tacitly or explicitly, a constraint, it limits the theoretical contents we could rationally countenance (the range of theories we *might* accept) and the theoretical contents that we do in fact accept. For instance, for any theory, there is some set of data that we wish the theory to account for. That set of data acts as a constraint on the theory. We expect our theory to be *adequate* to the data — that is, to be roughly consistent with it. As a rule, we don't accept theses that are inconsistent with our data unless there are overriding factors, like the theses are indispensable products of an otherwise empirically successful theory. We also expect our theory to *explain* the data, so our explanatory expectations also constrain the sorts of theoretical contents we entertain. All sorts of things can constrain our theories aside from the demands for empirical adequacy and explanatory power. We might require our theories to be, *inter alia*, internally consistent, unificatory, consilient with the broader theoretical landscape, or virtuous (simple, beautiful, elegant, convenient, familiar, fruitful). Sometimes we require them to be consilient with political or ideological goals. So for any given theory, the constraints we place on its content can depend on our local epistemic norms and pragmatic goals, as well as on features of our broader intellectual and socio-political context.

Theoretical constraints fall on a spectrum of strength. If our data set is small, our adequacy to the data constraint will not be particularly strong, since a good deal of theoretical content is consistent with a relatively small data set. The larger our data set, the stronger the constraint, since a relatively larger data set is consistent with relatively less theoretical content. The strength of our

adequacy constraint will also depend on our auxiliary hypotheses. Without any auxiliaries, most theories have no empirical consequences. With the right auxiliaries, just about any theory can be made consistent with any data. So if we start gerrymandering our theories to include ad-hoc auxiliaries, our adequacy constraint will be weak. If we limit the sorts of auxiliaries we countenance, we strengthen the constraint considerably.

Call a theoretical constraint *robust* when it falls on the high end of the spectrum — that is, when it permits relatively little theoretical content into a theory and rules out relatively much. A theory is *robustly constrained* when it has some robust constraints on its content. Call a constraint *permissive* when the opposite holds — that is, when it permits much content and rules out little. Note that these notions are relative, not categorical. The more robust (or less permissive) a constraint, the less content it permits into our theory; the more permissive (or less robust), the more it permits. A robust theoretical constraint is like a selective bouncer at a club, who lets in only the few people who meet his strict criteria for inclusion. It is good that he does so, because a club that is too inclusive is not a very good club. Likewise, I will argue, a theory that permits too much content is not a very good theory.

By *theoretical content*, I just mean *what the theory says*. There are many ways of carving up the content of a theory. We might talk about the number of theses in the theory — but this raises the question of how to individuate theses. We might talk about the number of *substantively* distinct claims — but this raises the question of what makes claims substantively distinct from one another. We might try to distill some minimal number of indispensable objects of quantification and predication and then count up how many things are predicated of how many objects. The issue is extremely thorny. I will remain neutral about the carving up of theoretical content and just stick

with the simple, intuitive gloss that I started with — the content of a theory is just what the theory says.

I will not argue for the blanket claim that the more constrained the theory, the better. On the contrary, there is such a thing as too much constraint. Theories that are *too* constrained cannot say much at all. Think, for instance, of Descartes' *Meditations*, in which Descartes' metaphysical theory is constrained by so demanding a standard of justification that it permits the theory few, if any theses. My point will be, rather, that there is an ideal range on the spectrum between *anything goes* and *nothing goes* — and that range is closer to the latter extreme than to the former.

Note that the notion of robust constraint that I have introduced is vague. However, we can identify the extreme ends of the spectrum of constraint based on whether theories say too much or too little to even have a chance of meeting their explanatory aims. Constraints are too robust or too permissive relative to certain aims that I take to be integral to the activity of theorizing.<sup>24</sup> Theories by their very nature are explanatory endeavours. A theory that is too robustly constrained can hardly say anything and as such doesn't have a chance at meeting its explanatory aims. That is, it is hardly a theory at all. Likewise, a theory whose constraints are too permissive cannot meet its explanatory aims either, but for different reasons. A permissive theory simply entertains too much to be explanatory — the thing we aim to understand and account for gets lost in a flood of irrelevant details. So it is in virtue of the kind of thing that theories are that some constraints are too robust or too permissive.

Robustness is not the only feature of theoretical constraints we should care about. Being permissive is not sufficient reason to abandon a constraint. Whether we should do so depends on

---

<sup>24</sup> I take the question of what a theory is to be distinct from the question of what the ultimate epistemic goods are.

whether there is a good rationale for the constraint. Suppose that we were to build a theory from scratch and make the demand for internal consistency our only theoretical constraint. It wouldn't matter what the theory was supposed to be about, since we wouldn't hold it accountable to any data and we wouldn't place any explanatory demands on it. The theoretical options would be limitless — that is, any number of different theories could satisfy the constraint. The internal consistency constraint rules out infinitely many inconsistent theories, but permits infinitely many consistent ones. So it is a permissive constraint. That *doesn't* mean we should give up on internal consistency! If the internal consistency constraint is a good one, it is so quite apart from how robust or permissive it is, but rather because contradictions are bad or problematic in some way.<sup>25</sup> So again, I don't wish to argue that we should abandon constraints solely on the grounds that they are permissive.

Likewise, being robust is not sufficient reason to implement a theoretical constraint. For instance, suppose we are developing a metaphysical theory and we adopt a theoretical constraint that allows us to countenance only theses that demonstrate the existence of God. The constraint would be robust, since little, if anything, would make it into the theory. But there is no epistemic reason to adopt such a constraint (of course, there might be non-epistemic reasons). This goes to show that our selection of theoretical constraints should be well-motivated. I don't wish to argue that we should adopt constraints solely on the grounds that they are robust.

What I wish to argue for, rather, is that we should, as a rule, robustly constrain theories. There may be good reasons to accept permissive constraints, such as the internal consistency constraint, but we ought to supplement those permissive constraints with other, more robust

---

<sup>25</sup>I have made this sentence conditional because I don't mean to take a stand on dialetheism here. We might not wish to take it for granted that contradictions are always bad or problematic (see Priest 1998).



constraints on theoretical content. Suppose our club bouncer is singularly dedicated to preventing weapon-carrying persons from entering the club. Supposing our club is in, say, an affluent part of Manhattan, that criterion for entry will be pretty permissive. But the criterion is nevertheless a good one! This just goes to show that we don't *only* care about the exclusivity of our club; among other things, we also care about how safe it is. But to the extent that we do care about exclusivity, we should put additional bouncers with more selective criteria at the door.<sup>26</sup> In the case of theories, I don't wish to argue that we should care *only* about how constrained our theories are. I wish to argue simply that we should care about it *inter alia*. When a theory is constrained only by a small number of permissive constraints, we should hold the theory accountable to additional, more robust ones.

Here's a broadly naturalistic argument in favour of robustly constraining theories:

**Argument from science:**

P1) Science is robustly constrained.

P2) Being robustly constrained has benefited science.

---

C) Other forms of inquiry should aim to be robustly constrained.

As I will argue in Chapter 3, science is robustly constrained (P1) because its experimental methods frequently lead it to disconfirm claims. One of the constraints that we frequently place on scientific theories is that they be accountable to the phenomena in various ways. Our scientific theories usually have to be adequate to a set of data that grows, on the whole, over time (but not always continuously, as in cases where auxiliary hypotheses allow us to disregard aberrant data; see McLeod and Parsons 2013) and they frequently have to explain and novelly predict it. The novel

---

<sup>26</sup> The analogy breaks down here, in that club bouncers can and do use multiple criteria to judge who gets in. But let's imagine that each bouncer can only use a single criterion.

prediction constraint is an especially robust constraint, since novel predictions are so hard to get — relatively few putative theoretical contents enable them.

Requiring of science that it be robustly constrained — that it be both empirically adequate and empirically successful — has benefited it in various ways (P2). Science receives a good deal of its epistemic distinction from its empirical success. In fact, one of the main arguments for scientific realism — the no-miracles argument — infers from the great empirical success of science its approximate truth. But to say that science is empirically successful is to say that it has been successful in generating theoretical content that respects the rigorous constraints partly constitutive of its methods. So the epistemic distinction that we might assign to science owes in part to its being robustly constrained. Being able to successfully explain and predict phenomena also has certain pragmatic benefits: it facilitates successful manipulation and intervention. For instance, the understanding and predictive capability science gives us allows for all sorts of technological and medical payoffs. So being robustly constrained has had both epistemic and pragmatic payoffs for science.

The argument has some pull, but its conclusion, C, that non-scientific forms of inquiry should be robustly constrained, follows only on the assumption that what benefits one form of inquiry will similarly benefit other forms of inquiry. As a general claim, that would be hard to substantiate. If we restrict C to a claim about *metaphysics* in particular, our assumption would have to be only that *metaphysics* would benefit from robust constraint in a manner similar to science. We might have less difficulty supporting that second assumption, since metaphysics and science share at least some basic aims, including the aim of describing the underlying nature of reality. Still, it is not clear how to go about defending it. One way might be by demonstration — by examining robustly constrained metaphysics and surveying any epistemic and pragmatic successes

owing to that constraint. That would be an illuminating project, but it is outside the scope of this dissertation. So I will not defend the argument from science above. Instead, I wish to articulate more general reasons to think that robust constraint is a sound epistemic policy. That is, I wish to motivate some epistemic principles having to do with robust constraint. If any of the principles is true, then we have some motivation for something like C above — that is, for thinking that as a rule, we should robustly constrain our theories.

Here are a number of candidate principles:

- Weak Constraint Principle:** *Ceteris paribus*, theories that are robustly constrained are epistemically preferable to those that are not.
- Moderate Constraint Principle:** Theories that are robustly constrained are epistemically preferable to those that are not.
- Strong Constraint Principle:** To be epistemically adequate, a theory must be robustly constrained.

The weak principle makes robust constraint preferable *ceteris paribus*; the moderate principle makes it preferable full-stop; the strong principle makes it a necessary condition of what I call *epistemic adequacy*. For a theory to be epistemically adequate is for it to be permissible for us to accept it for epistemic rather than pragmatic or heuristic purposes, where we are remaining neutral on the exact normative force of *permissible*, and where *accepting* is to be understood in L. Jonathan Cohen's sense:

to accept that *p* is to have or adopt a policy of deeming, positing, or postulating that *p* — i.e. of including that proposition or rule among one's premisses for deciding what to do or think in a particular context, whether or not one feels it to be true that *p*. (1992, 4)

When I say that a theory is *epistemically inadequate*, I mean that from an epistemic (as opposed to pragmatic) standpoint, we shouldn't accept it, in Cohen's sense. At any rate, of the three principles, I favour the strong one, but there are compelling arguments in favour of all three. I will survey those arguments in section III below. Again, if *any* of the three principles is correct, then my larger conclusion follows. That is, supposing that we should pursue what is epistemically preferable or preserve epistemic adequacy, then in general, *we should robustly constraint our theories*.

In this section, I introduced the notion of theoretical constraint – and of robust constraint, in particular. The main purpose of the section was to introduce three candidate epistemic principles revolving around the notion of robust constraint: one weak, one moderate, and one strong. The principles claim that robust constraint is either an epistemic plus or an epistemic requirement. I will motivate these principles in Section III below. For my purposes, the importance of the principles is that if any one of them is true, then we have some motivation for my broader conclusion: that as a rule, we should robustly constrain our theories. The next section will introduce the second set of epistemic principles that I aim to motivate in this chapter.

## **II.2 Epistemic Warrant**

The second set of epistemic principles concerns epistemic warrant. By *epistemic warrant* (or, simply, *warrant*) I mean justification. I will use these terms interchangeably. The nature and features of justification are a matter of much controversy — controversy too deep to treat at any length here. In short, there are two main classes of theory about the nature of justification. According to *internalist* theories, being justified is a function of having the right internal states — being conscious of the relevant evidence, for instance (see Chisholm 1989 and Conee and Feldman 2001, 2004). *Externalist* theories hold that it is a function, at least partly, of things external to the

agent, such as the worldly reliability of one's belief-forming processes (see Alston 1989 and Goldman 1979, 1999). The internalist concerns herself mainly with particular epistemic circumstances and the externalist mainly with epistemic policy — much like the act utilitarian concerns herself with individual moral deliberations and the rule utilitarian with moral policies. Both the individual cases and the policies are important in my view. However, I won't avow any theory of epistemic warrant here.

Let me instead stipulate a number of things. On my understanding of the term, being warranted or not is feature of propositions and theories. It admits of degrees — a proposition or theory can have greater or lesser epistemic warrant. It is also relational. Propositions or theories are warranted in relation to certain epistemic agents with certain evidence. Internalists and externalists can agree about the general importance of evidence. Even if justification requires belief-forming processes that are reliable in some world-relative way, those processes will in many cases involve the collection of evidence. As an epistemic policy, collecting evidence for the propositions one believes or accepts is good. Set aside whether reliable processes of belief formation *must* involve consciously available evidence. For my purposes, I'll understand epistemic warrant as follows:

A theory or proposition has some degree of **epistemic warrant** for some epistemic agent at  $t$  if she has evidence for it at  $t$ .

To be evidence of  $p$  is to be something that supports  $p$ . By that, I mean that the truth of the evidence makes  $p$ 's likelihood greater than it would be otherwise. I am open to  $p$ 's being explanatory being evidence in its favour. I remain neutral about what kinds of things constitute evidence (propositions, events, mental states, and so on). The greater the evidence quantitatively and qualitatively, the greater the epistemic warrant. *Prima facie*, this sounds like an internalist

understanding of warrant. However, I don't claim that evidence is necessary for warrant (the definition above does not say *only if*); I claim only that it is sufficient. Moreover, since I haven't made evidence a necessary condition of warrant, I haven't committed myself to evidentialism. So we are still on relatively neutral ground.

Having laid the conceptual groundwork and flagged my assumptions, I'm now in a position to present the next set of epistemic principles. Where *warrant* is understood in the above sense:

**Weak Warrant Principle:** *Ceteris paribus*, claims that are sufficiently warranted are epistemically preferable to those that are not.

**Moderate Warrant Principle:** Claims that are sufficiently warranted are epistemically preferable to those that are not.

**Strong Warrant Principle:** To be epistemically adequate, theories must have sufficient warrant.

The weak warrant principle makes sufficient warrant *ceteris paribus* epistemically preferable; the moderate principle makes it epistemically preferable full-stop; and the strong principle makes it a condition of epistemic adequacy, where *epistemic adequacy* is to be understood as above. Again, I think the strongest version of the principle holds. But so long as just one of them holds, the second of my broader conclusions is secure. Whichever principle we prefer, on the assumption that it is good to be epistemically preferable or epistemically adequate, then in general, we should seek adequate warrant for our theories.

The warrant principles are not vague. That is because the notion of 'sufficient' or 'adequate' warrant can be made precise. By *sufficient* or *adequate epistemic warrant*, I mean that

an ideal Bayesian reasoner would assign a greater-than-chance subjective probability to the theory given the evidence. Where T stands for the relevant theory, Bayes' theory tells us:

$$P(T|E) = P(T) \times \frac{P(E|T)}{P(E)}$$

If the ideal reasoner assigns her subjective probabilities such that  $P(T|E) > 0.5$ , then she has sufficient epistemic warrant. Note that this is a low bar for adequacy — we might call theories to which the ideal Bayesian assigns a just-greater-than-chance credence *minimally* adequate. Clearly, the greater the ideal credence, the better.

I have laid out a set of constraint principles and a set of warrant principles, from which the following norms follow, respectively: 1) theories should be robustly constrained and 2) theories should be adequately warranted. In Chapter 3, those norms will figure into my evaluation of free range and scientifically responsible metaphysics. But presently, it remains to be seen why theories should be held to them. Why should we robustly constraint and adequately justify our theories? The answer is that we have reason to accept at least one of the constraint principles and at least one of the warrant principles. In the following section, I will provide a number of arguments in favour of the various principles.

### **III. Defending the Epistemic Principles**

#### **III.1.1 Constraint Principles**

Robust theoretical constraint has a number of potential benefits. First, by limiting the theoretical contents we countenance or accept, constraints make our theories more likely to be true. This is a purely statistical matter, relating to *how much* the theory claims. The argument goes as follows.

**Argument from statistical likeliness:**

- P1) *Ceteris paribus*, sparser theories are statistically likelier to be true than (non-nested) rivals with more theoretical content.  
P2) *Ceteris paribus*, statistically likelier theories are epistemically preferable.  
P3) Robustly constrained theories are relatively sparse.  
P4) Robustly constrained theories are statistically likelier to be true than (non-nested) rivals with more theoretical content.

---

C) *Ceteris paribus*, robustly constrained theories are epistemically preferable.

By “sparser” theories, I do not mean logically weaker theories; I mean simpler theories or theories with fewer adjustable parameters.<sup>27</sup> The idea behind P1 is that *ceteris paribus*, the more adjustable parameters there are in the content of your theory, the less statistically likely your theory is to be true. Now, let me acknowledge that Popper famously argued the opposite — that simpler theories are less statistically likely. Take the following example from Sober:

(LIN) There exist numbers  $a_0$  and  $a_1$  such that  $y = a_0 + a_1x$ .

(PAR) There exist numbers  $a_0$ ,  $a_1$ , and  $a_2$  such that  $y = a_0 + a_1x + a_2x^2$ . (2015, 89)

Insofar as LIN contains fewer adjustable parameters than PAR (it lacks  $a_2$ ), LIN is simpler than PAR. But, Popper’s point is that LIN cannot be more probable than PAR. That is because LIN is equivalent to  $(PAR) \wedge (a_2 = 0)$ , and a conjunction is never more probable than its individual conjuncts. So we have a case where the simpler of two theories is less statistically likely.<sup>28</sup> However, Sober points out that “when Popper talks about examples like LIN and PAR, he takes these models to be *mutually compatible*; in fact, one of them logically entails the other. Statisticians describe this relationship by saying that LIN is *nested* inside of PAR” – and anyone claiming that the simpler theory is statistically likelier assumes that the rival models are not nested in that manner

---

<sup>27</sup> Note that these can come apart (see Sober 2015, 93).

<sup>28</sup> I thank Peter Godfrey-Smith for pressing me to consider this point.



(2015, 92). So that is why P1 above contains the parenthetical qualification that the rival theories must not be nested.

P1 finds support in statistical arguments for Ockham's Razor. Harold Jeffreys (1939) famously claimed that we should assign relatively higher prior probabilities to relatively simpler claims. But this won't do, because without some epistemically-motivated reason to assign priors in that fashion, it just begs the question at issue (see Kelly 2004 and Sober 2015). A more compelling reason to think that simpler theories are likelier is that a Bayesian form of Ockham's Razor emerges naturally in Bayesian reasoning (Nichols et. al 2016, 534 fn. 1).<sup>29</sup> The view is that, since "models that contain fewer adjustable parameters postulate a narrower range of possibilities over which probabilities must be distributed", models with fewer adjustable parameters are likelier (Sober 2015, 125). One particular version of Bayesian Ockham's Razor is the *size principle*, which has emerged out of cognitive science (see Nichols et al. 2016, Perfors et al. 2011 and Tenenbaum and Griffiths 2001). According to the size principle, "smaller hypotheses that are consistent with the data... are significantly preferred to larger hypotheses... and this advantage increases exponentially with each new data point", where the size of the hypothesis has to do with the size of any sets it quantifies over (Nichols et. al. 2016, 534).<sup>30</sup> So the idea is that the less your claim quantifies over, the more statistically likely it is. Parsimony is also relevant in non-Bayesian contexts where we prefer models that are predictively accurate, since it helps us assess the predictive accuracy of models (Sober 2015, 130-135; see also Forster and Sober 1994). For instance, parsimony is relevant to assessments of predictive success in machine learning theory (Sober 2015, 140 fn. 61). The point is that on several model selection criteria, parsimony is

---

<sup>29</sup> On Bayesian Ockham's Razor, see also Berger (1985), Box and Tao (1973), Henderson et al. (2010), Jefferys and Berger (1992a), MacKay (2003), and Sober (2015).

<sup>30</sup> I thank Eric Mandelbaum for drawing this to my attention.

mathematically relevant to estimating the model's predictive accuracy and likelihood (Sober 2015, 141). In other words, "Ockham's razor is alive and well in statistics" (Sober 2015, 141). If so, then sparseness conduces to likeliness (P1).

That means that, if the greater statistical likeliness of one's theory is an epistemic good, then (P2) *ceteris paribus*, we should prefer the sparser theory. Note that P2 hinges on greater statistical likeliness having epistemic value, which I won't defend here. Again, I do not need to defend specific epistemic values given the conditional position I am defending. The *ceteris paribus* clause is crucial, because there are a good deal of other features of theories that we care to preserve, the importance of which trumps that of statistical likeliness. For instance, we have certain descriptive and (as I have already said) explanatory aims in formulating theories. We want our theories to be informative, contribute to understanding, describe their subject-matter in a certain level of detail, and be as complete as possible. On the whole, it is better to have a descriptive and explanatory theory that is only somewhat likely than to have a theory that is not particularly descriptive or explanatory, but that is a good deal likelier. So we *don't* want a principle according to which the less our theory says, the better. That would have the silly consequence that we should prefer less informative theories to more informative ones. We want, rather, a principle that tells us to prefer the sparser of two theories that are equally satisfactory in other important respects.

Robustly constrained theories are relatively sparse (P3) because being robustly constrained is a matter of there being relatively fewer degrees of freedom with respect to the theoretical content that we countenance or accept. Being robustly constrained just is a matter of limiting the theoretical content we countenance or accept — that is, of making our theories sparse relative to less constrained actual and potential theoretical alternatives. It follows from the claim that sparser theories are statistically likelier (P1) and the claim that robustly constrained theories are relatively

sparse (P3) that robustly constrained theories are statistically likelier (P4). Finally, having established P1-P4, it follows that robustly constrained theories are *ceteris paribus* preferable (C). So we have a probabilistic argument for our weak constraint principle.

We may motivate a second argument for the weak constraint principle by noting that robust constraint is conducive to agreement, which we might also take to be an epistemic good.

**Argument from agreement:**

P1) *Ceteris paribus*, relatively sparse theories are more conducive to agreement.

P2) *Ceteris paribus*, theories that conduce more to agreement are epistemically preferable.

P3) Robustly constrained theories are relatively sparse.

---

C) *Ceteris paribus*, robustly constrained theories are epistemically preferable.

Just as a sparser theory is more statistically likely, a sparser theory is also, *ceteris paribus*, more conducive to agreement (P1). That is because, the fewer substantive commitments our theory makes, the less there is to disagree about. In fact, I will note in section I.3.5 of Chapter 3 that, notwithstanding the presence of disagreement across most epistemic contexts, deep, persistent, nebulous disagreement can be a symptom of poor constraint. The less well-constrained a theory is, the greater and more pervasive disagreement will be, and the wider the range of theoretical alternatives potentially on the table. Religion illustrates the point well, since humans have posited such a rich variety of religious theories throughout history and since our evidence so dramatically underdetermines those theories. The point is that robust constraint and disagreement inversely correlate. The *ceteris paribus* clause in P1 is meant to catch unusual cases in which a fuller, more descriptive theory commits only to truisms, while a less descriptive one makes controversial claims (for instance, a theory that commits to hundreds of simple mathematical equations versus a theory that says only that abortion is morally wrong).

The claim that whatever conduces to greater agreement is *ceteris paribus* preferable (P2) hinges on an assumption that agreement is an epistemic good, which, again, I won't argue for here. If agreement is an epistemic good, then P2 follows. The *ceteris paribus* clause is meant to signal that, as in the case of statistical likeliness, agreement can be trumped by other factors, given our descriptive and explanatory aims. A fully descriptive and explanatory theory that generates substantial disagreement is preferable to a theory that we can agree on, but that contributes little to our understanding. I established the claim that robustly constrained theories are relatively sparse (P3) in my explanation of the argument from statistical likeliness above, and the conclusion that robustly constrained theories are *ceteris paribus* epistemically preferable follows. So if we take agreement to be an epistemic good, then we have further reason to accept the weak constraint principle.

In sum, when theoretical constraint is robust and therefore permits a narrow range of theoretical content, there are two potentially valuable outcomes: greater statistical likeliness of our theory and greater consensus about that theory. If either of those outcomes is epistemically valuable, it follows that we should prefer robustly constrained theories, other things being equal.

We have already seen an argument built around the claim that sparser theories are statistically likelier. A related claim is that the more theoretical contents we rule out and the fewer we permit, the likelier we are to avoid falsity. The argument goes as follows.

**Argument from falsehood avoidance:**

- P1) The more content we exclude from a theory and the less we permit, the more likely we are to avoid falsity.
- P2) *Ceteris paribus*, theories that are more likely to avoid falsity are epistemically preferable.
- P3) Robustly constrained theories exclude relatively many putative theoretical contents and permit relatively few.

C) *Ceteris paribus*, robustly constrained theories are epistemically preferable.

The idea is that the less you say, the less likely you are to say something false. This is so for purely probabilistic reasons – namely that  $P(p) > P(p \wedge q)$ .<sup>31</sup> Now, it's certainly true that in some toy cases, we are as a matter of fact better off saying more than less. For instance, suppose that  $p$  is false and  $q$  is true. If I say only  $p$ , I've said something false. If I say  $p$  and I say  $q$  (treat them as independent claims, not a conjunction), then I've said something false and something true. So on the whole, I'm better off having said more than less.<sup>32</sup> But here we are building in assumptions about the truth-values of  $p$  and  $q$  and stipulating that when I say less, I say the false thing. The point stands that from a statistical standpoint, where we don't know the truth-values of  $p$  and  $q$ , the less you say, the less likely you are to say something false. For anyone who takes falsehood avoidance to be an epistemic good, making our theories sparser is good.

The claim that avoiding falsity is *ceteris paribus* epistemically preferable (P2) assumes that likeliness to avoid falsity is an epistemic good. As in the previous arguments, our *ceteris paribus* clause does some needed work. It signals that making falsity less likely is an important epistemic aim only *inter alia*. We would of course prefer a theory that tells us something about our subject matter and risks some falsity over a theory that tells us next-to-nothing and does not risk falsity.

The claim that robustly constrained theories exclude relatively many putative theoretical contents (P3) is just built into the notion of robust constraint. By definition, part of what makes a constraint robust is that it excludes a broad range of putative theoretical contents from consideration. It follows that robustly constrained theories are *ceteris paribus* epistemically preferable (C).

---

<sup>31</sup> I thank Aaron Ancell and Dmitri Gallow for helpful discussion of this point.

<sup>32</sup> I thank Michael Devitt for the example.

We have seen in the discussion of the arguments so far the importance of explanatory and descriptive aims. Every theory aims to adequately account for some subject-matter. That is, there is a set of facts with which the theory aims to acquaint us. Assuming that not every claim about some given subject-matter is true, then the facts we aim to discover are to some degree limited. A poor method of discovering those facts would be to countenance theses at random, stumbling around in possibility space, hoping that by sheer luck we stumble upon some facts. The method is poor because it is inefficient given the size of possibility space.<sup>33</sup> A more efficient method would better *target* the desired facts. While the arguments we have seen so far motivate only the weak constraint principle, this last observation motivates the moderate constraint principle.

**Argument from methodological efficiency:**

- P1) Methods of theory construction that more effectively target the relevant facts are epistemically preferable.
- P2) Ensuring that our theories are robustly constrained allows them to more effectively target the relevant facts.

---

C) It is epistemically preferable to robustly constrain our theories.

Robustly constraining our theories makes our investigation more efficient. Unlike the previous arguments, we do not need a *ceteris paribus* clause. If efficiency is epistemically preferable (P1), then it is so as a rule rather than *ceteris paribus*, because there is little to be gained by targeting facts *less* effectively — no obvious epistemic benefit, the importance of which would trump the importance of targeting the desired facts well. That is why I take the argument to establish the

---

<sup>33</sup> One *pragmatic* reason why we care about efficiency is that, as I noted above, we are finite inquirers with finite resources. It is in our interests for our inquiries to proceed efficiently, since we have limited time.

moderate constraint principle rather than the weak one (recall that what distinguishes the weak from the strong is that the weak one has a *ceteris paribus* clause).

Methods that target the relevant facts effectively are epistemically preferable (P1) assuming that some kind of methodological expediency and systematicity is an epistemic good. Further, targeting facts helps us to discover them more readily. That is, effective targeting is truth-conducive. So if true belief is an epistemic good, then effective targeting is good insofar as it conduces to true belief. Moreover, if we do not target our desired facts effectively, we risk countenancing a good deal of putative theoretical contents that are irrelevant to our subject-matter. So targeting the desired facts effectively helps maximize the relevance of the claims we entertain to the subject-matter we aim to account for. In that respect, it makes the theory better able to do what it sets out to do — namely to fulfill its explanatory aims — since relevance is a key component of satisfying explanations. So if maximizing relevance is an epistemic good, then effective targeting is good insofar as it conduces to it. In sum, effective targeting enables a number of things that we might consider epistemic goods: methodological expediency and systematicity, truth-conduciveness, and the maximization of relevance.

Robust constraint is conducive to effective targeting (P2) because part of the job of a constraint is to help the theory target the desired facts. When theoretical constraints screen off large amounts of putative theoretical content, they narrow the range of theoretical contents that the theoretician could countenance or accept into her theory. Theoretical constraints direct the theoretician's attention away from large swaths of possibility space, toward others. The more robustly constrained our theories, the more narrowly they target the desired facts. Note, however, that the screening off function doesn't *suffice* for effective targeting – it just conduces to it. So if targeting the desired facts is like directing a flashlight toward a specific region, we want the beam

of light to be a) suitably narrow *and* b) pointing in the right direction. The fact that robust constraint rules out a bunch of putative theoretical contents gets you just the first thing but not the second; the additional requirement that our constraints be well-motivated would hopefully get us the second. At any rate, since robust constraint helps with targeting, it follows that robustly constrained theories are epistemically preferable as a rule (C). So we have some motivation for the moderate constraint principle.<sup>34</sup>

We have looked at a number of arguments. The arguments from statistical likeliness, agreement, and falsehood avoidance supported a weak constraint principle, according to which we should prefer robustly constrained theories *ceteris paribus*. This last argument from methodological efficiency supported the moderate constraint principle, according to which we should prefer robustly constrained theories as a rule. In fact, I think a stronger version of the argument from methodological efficiency is defensible, which I will defend presently.

**Argument from methodological inefficiency:**

P1) Methods of theory construction that fail to effectively target the relevant facts produce epistemically inadequate theories.

P2) Failing to robustly constrain a theory is a failure to effectively target the relevant facts.

---

C) Failing to robustly constrain a theory makes it epistemically inadequate.

As an epistemic policy, we should not accept theories produced by unreliable methods — that is, by methods that are not reliably truth-tracking or that do not reliably produce justification. That is

---

<sup>34</sup> Compare Kelly (2004) who argues, similarly, that simplicity helps us to get at the truth *more efficiently*, since opting for simpler models will mean fewer theory changes en-route to the truth. Kelly explains:

disregarding Ockham's advice opens you to a needless, extra U-turn or reversal in opinion prior to all the reversals that even the best of methods would have to perform if the same answer were true. So you ought to heed Ockham's advice. Simplicity doesn't indicate the truth, but it minimizes reversals along the way. (2004, 492)



because the more unreliable the method, the riskier it is epistemically to accept the theoretical products of that method. A failure to effectively target facts is a failure to be reliably truth-tracking. So a method that fails to effectively target facts will produce epistemically inadequate theories (P1).

To see why the failure to robustly constrain one's theory is a failure to effectively target the relevant facts (P2), think about what a theoretical constraint does. It limits what our theory can say. The fewer the constraints, the more the theory can say. If you do not constrain the content of your theory, then your theory can say anything whatsoever. If you do not *robustly* constraint it, then the range of claims that your theory could commit to is enormous. Where that is the case, the theory is far less likely to be true or approximately true on the whole. If it does get at the truth, it does so by luck and not in an efficient, directed manner. So to target the relevant facts effectively, one needs to robustly constrain one's theoretical contents. It follows that robust constraint is a necessary condition of epistemic adequacy (C). So there is compelling reason to accept the strong constraint principle.

In sum, there are a number of arguments in favour of the various constraint principles. Considerations of statistical likeliness, agreement, and falsehood avoidance supported the weak constraint principle. Considerations of methodological efficiency and inefficiency supported the moderate and strong constraint principles, respectively. Given a compelling argument for the strongest constraint principle, that is the principle I am inclined to accept. But at any rate, so long as one of the principles is true, then as a rule, we should robustly constrain our theories.

### **III.1.1 Robust Constraint and Simplicity**

The talk of sparsity in some of the arguments above might lead one to wonder whether the constraint principles I have articulated are really just demands for simplicity. The two are related.

Adhering to the constraint principles produces relatively simpler theories. However, robust constraint and simplicity are not the very same thing. That is because a theory can be simple and not robustly constrained and, likewise, it can be robustly constrained and not simple. We will see an example of the former in Chapter 3. In brief, the example is of a theory according to which the world contains a single apple. That theory is simple, but to the extent that we explicitly require of the theory only that it be logically consistent, it is not robustly constrained. As for a robustly constrained theory that is not simple, imagine some mathematical or logical system, highly constrained by axioms, in which the mathematician or logician introduces some unnecessary — that is, superfluous — posit for the purposes of convenience. The system would be robustly constrained, but not completely simple. Robust constraint and simplicity come apart, so the demand for the one is not the demand for the other.

### **III.2 Warrant Principles**

Since a good deal more has been written about epistemic warrant than about theoretical constraint, defending the warrant principles is somewhat easier, in that a lot of the work of arguing for the epistemic value of justification has already been done by others. As we have seen, the nature of justification and its relation to knowledge is a matter of much controversy. But whether we understand justification in internalist or externalist terms, whether or not we think justification is necessary for knowledge, and regardless of what we take to be the *ultimate* epistemic aims or values, surely justification is *an* epistemic aim or value. That is, surely it is not a matter of much controversy that having evidence for our claims is in general good from an epistemic point of view. There are several reasons why this is so.

First, it is a near-consensus view that justification aims at truth, or is a means to true belief.<sup>35</sup> Given this point of relative agreement, disagreements about the nature of justification can be understood as disagreements about the connection between justification and truth (Lehrer and Cohen 1983). At any rate, on this near-consensus view, one justifies one's beliefs in order to make it more likely that they are true. If so, and if true belief is epistemically valuable, then justification is epistemically valuable to the extent that it promotes true belief.

Many of those who explicitly tie justification to the pursuit of truth also tie it explicitly to the avoidance of falsehood (e.g. Alston 1989, BonJour 1985, Foley 1987, Valid 2003, Weiner 2005). On that view, justifying one's beliefs makes them less likely to be false. If so, and if the avoidance of falsehood is epistemically valuable, then justification is epistemically valuable to the extent that it aids falsehood avoidance.

Further, having justification is a necessary condition of having justified true belief. As we have seen, some argue that *justified* true belief is epistemically valuable over and above true belief (DePaul 2001, Feldman 2002, Smithies 2012). If so, then acquiring justification is epistemically valuable to the extent that it is necessary for justified true belief.

Moreover, justification may have epistemic value based on its connection to rationality. Some argue that justification and rationality are intimately bound up. Langsam (2008) argues that debates about the nature of justification are in fact debates about the nature of rationality. The precise nature of the connection between the two is debatable: Long (2010) and Swinburne (2001) think they are equivalent, Betz (2013) takes justification to be a condition of rationality, and

---

<sup>35</sup> See Alston (1985), (1986), (1989); Audi (1988); Betz (2013); BonJour (1985); Chisholm (1980); Foley (1987), (1993); Giere (1989); Kitcher (1992); Laudan (1990a); Lehrer and Cohen (1983); Nozick (1993); Sosa (1980); Steglich-Petersen (2013); Swinburne (2001); Vahid (2003); Velleman (2000); and Weiner (2005). See Cruz and Pollock (2004) for a counterpoint and Graham (2011) in response.

Langsam (2008) takes justification to be both necessary and sufficient for rationality. Regardless of the precise relation between justification and rationality, the two are intimately related. For that reason, if rationality is an epistemic good, then whatever epistemic value attaches to rationality likely attaches to justification, too.

So there are several candidate reasons to assign epistemic value to justification or warrant. Those reasons motivate at least the weak warrant principle, according to which warranted claims are *ceteris paribus* epistemically preferable to unwarranted ones. It motivates only the weak principle insofar as *having epistemic value* could mean having *defeasible* epistemic value — that is, epistemic value that is trumped by greater epistemic goods.

But what sorts of epistemic goods could trump the value of warrant? There might be cases where the expedience of my belief formation takes precedence over warrant. Evolution might favour quick judgments in certain sorts of cases over accurate or epistemically warranted ones. For instance, from an evolutionary perspective it might be more advantageous to me that I quickly (though typically falsely) judge little shadowy things to be spiders, rather than investigating every little shadowy thing before I judge it to be a spider. But the greater evolutionary value of expedient belief formation is not an *epistemic* value; it is a pragmatic one. The question is whether any *epistemic* values trump the value of warrant.

We might think that *true belief* has greater epistemic value than warrant. And sometimes, true belief might come apart from warranted belief. Evidence can be misleading, after all. Suppose that some epistemic agent must believe either *T* or *W*, but not both, where *T* is true but unwarranted and *W* is warranted but false. Is it epistemically better for her to believe *T* or *W*? There are arguments to be made on both sides. On this occasion it may be, objectively speaking, better for her to believe *T*; as a matter of epistemic policy it may be better for her to believe *W*. But at any

rate, our epistemic agent is in limited epistemic circumstances. She must choose not between true belief and evidence, but between *what she believes to be true* and what she has evidence for. And if she believes  $T$  to be true, she should probably have evidence for  $T$ . But we have already said that  $T$  is unwarranted. So it is not clear what epistemic reason she has to believe  $T$  and not  $W$ . All this goes to show that while the epistemic value of true belief might trump that of warrant objectively, it doesn't trump it in practice, among actual epistemic agents in limited epistemic circumstances. Absent some other candidate epistemic good, the value of which might trump the value of warrant, we don't need the *ceteris paribus* clause attached to the weak warrant principle. Therefore, the reasons I surveyed above for thinking that warrant is epistemically valuable support the moderate warrant principle, according to which warranted claims are epistemically preferable to unwarranted ones *tout court*.

Are there reasons to accept the strong warrant principle? Recall that the strong warrant principle makes warrant a necessary condition of epistemic adequacy. That is not to say that warrant is necessary for knowledge.<sup>36</sup> It is only to say that if it is epistemically permissible for us to accept  $p$ , then we have adequate evidence for  $p$ . By contraposition, if we don't have adequate evidence for  $p$ , it is not epistemically permissible for us to accept  $p$  — that is, we ought not accept it. The following argument establishes that conclusion.

**Argument from ideal Bayesian reason:**

- P1) For any proposition  $p$  for which I have insufficient warrant, my credence that  $(p|E) < 0.5$  if I am an ideal Bayesian reasoner.
- P2) As an epistemic policy, I ought not to accept propositions for which my credence is  $< 0.5$ .

---

<sup>36</sup> Recall that for  $T$  to be epistemically adequate is for it to be permissible for us to adopt a policy of claiming  $T$ . More is built into the notion of knowledge, including the requirement of truth. So epistemic adequacy isn't the very same thing as knowledge.

C) As an epistemic policy, I ought not to accept propositions for which I have insufficient warrant.

P1 is true by virtue of my definition of sufficient evidence. I ought not to accept propositions to which I assign a less-than-chance credence (P2), because where my credence in  $p$  is less than chance, I believe that  $p$  is more likely to be false than true, and it would be silly to make a policy of accepting propositions that I believe are more likely to be false than true. Therefore, given my Bayesian understanding of sufficient warrant, the strong warrant principle (C) follows.

I take this last argument to be compelling and therefore am inclined to accept the strong warrant principle. But as before, I need just one of the principles to be true in order to establish that we should adequately warrant our theories. Having surveyed a number of compelling arguments, we may safely accept the truth of at least one of them. We should, therefore, seek adequate warrant for our theories.

## **Conclusion**

I have argued that, just as there is an ethics of belief, there is an ‘ethics’ of theory building, insofar as we can intelligibly and fruitfully think about the norms that govern theory-building. Over the course of this chapter, I have established some of the contents of that ‘ethics’ of theory-building. In section II, I articulated two sets of principles of varying strength: one set to do with theoretical constraint and one to do with epistemic warrant. In section III, I presented a number of arguments in favour of the various principles. Ultimately, there is a compelling case for the strongest of both sets of principles. Regardless of whether we accept the stronger principles or the weaker ones, so long as one of the constraint principles is true and one of the warrant principles is true, it follows that as a rule, we should robustly constrain and adequately justify our theories. I

have shown that, if we take things like greater statistical likeliness, agreement, avoidance of falsity, methodological expediency and systematicity, truth-conduciveness, maximization of relevance, and rationality to be epistemically valuable, then we should accept at least one of the constraint principles and at least one of the warrant principles. While I do not have space to establish that those putative epistemic goods are actual epistemic goods, I think it reasonable to assume that some of them are. So we should robustly constrain and adequately justify our theories. Therefore, I will proceed in Chapter 3 to evaluate scientific and metaphysical theories based on whether they do so. Because I take the arguments for the strong principles to be sound, I will also consider any theory that fails to be robustly constrained or epistemically warranted to be epistemically inadequate.

## CHAPTER 3

### **On the Inadequacy of Free Range Metaphysics and the Appeal of Scientifically Responsible Metaphysics**

In this chapter, I present the two central theses of the dissertation. The first thesis rejects the view that free range metaphysics produces justified theories of reality. I argue that such metaphysical theorizing is epistemically inadequate because its constraints fail to meet the epistemic criteria established in Chapter 2 — that is, they are not sufficiently robust and their satisfaction does not secure adequate epistemic warrant. For those interested in limning the nature of the world, the second thesis prescribes naturalized or *scientifically responsible* metaphysics — metaphysics that is conscientiously engaged with the empirical data, theoretical insights, and practices of the current sciences. While the thesis that metaphysics should be naturalized or made accountable to science has a good deal of support and has given rise to scientific metaphysics — a burgeoning sub-discipline of metaphysics — more work needs to be done to justify and articulate the naturalization of metaphysics.<sup>37</sup> I argue that a scientific constraint on metaphysical theorizing is suitably robust and justificatory because scientific theses are systematically confirmed and disconfirmed.

---

<sup>37</sup> There is no difference in kind between what is commonly referred to as scientific metaphysics and what I call scientifically responsible metaphysics. I have simply chosen a different name, partly because I wish to introduce a term of art. But I also resist the name *scientific metaphysics* because it suggests, misleadingly, that its constituent metaphysical theories are themselves scientific — that they engage in the *methods* of science — when they need not be (§II.2 below). But what is now being called *scientific metaphysics* is exactly the sort of thing I am recommending here. Everything that qualifies as scientific metaphysics qualifies as scientifically responsible metaphysics.



## Part I: Free Range Metaphysics

### I.1 Definition

*Free range metaphysics* is metaphysics that science has only a nominal role in constraining. Academic metaphysics that is entirely *unconstrained* by science probably does not exist, since most educated people have undergone at least some basic level of science education. But science plays a minimal role in free range metaphysics. Its practices and institutional products — data, theories, books and journal articles — are not directly or explicitly appealed to. That is, free range metaphysics does not directly engage with science. Instead, it is constrained primarily by logical demands, such as the demand for consistency, by aesthetic demands, such as the demand for simplicity, and by psychological demands, such as the demands for intuitive plausibility and explanatory power.<sup>38</sup> I call these theories *free range* because the minimal constraints on their content permit the theories to go where they will. That is, any number of metaphysical theories could satisfy those constraints.

### I.2 Examples

Free range metaphysics is pervasive. Open a metaphysics anthology and you may find some scientifically-informed discussions of time, causation, natural laws, and so forth. But you will *certainly* find a number of discussions of personal and numerical identity, possible worlds, properties, substances, universals and other abstract entities, conducted independently of any

---

<sup>38</sup> One might wonder why I characterize simplicity as an aesthetic constraint, when in Chapter 2 I countenance some epistemic benefits of sparsity. In fact, I think that simplicity has both aesthetic and epistemic dimensions. The reader can take ‘aesthetic’ as shorthand for ‘partly aesthetic’.

scientific considerations.<sup>39</sup> These discussions often attempt to formulate the simplest and most explanatorily powerful theories consistent with the participants' intuitions. I will give three examples.<sup>40</sup>

The first example is van Inwagen's (1981) rejection of the doctrine of arbitrary undetached parts. According to that doctrine, for every material object occupying some spatial region at time  $t$ , a material object exists at *any* of the region's occupiable sub-regions at  $t$  (1981, 191). Van Inwagen argues against the doctrine by *reductio*, showing that in practice it entails violations of the transitivity of identity (1981, 195). So his rejection of the doctrine hinges on certain logical entailments being at odds with a feature of identity we take to be essential to it. The intuition that the feature is essential likely originates in conceptual understanding. Theories that rely only on this kind of argument and do not appeal to science constitute free range metaphysics. Let me emphasize that there is nothing intrinsically wrong with or inappropriate about appeal to logical

---

<sup>39</sup> I do not mean to suggest that the *topic* makes a metaphysical theory free range or otherwise. It is, rather, a matter of how the theory is constructed. Granted, there may be some metaphysical topics that belong to free range metaphysics necessarily — topics that science has no bearing on in principle. It would be difficult to *know* which topics science cannot speak to in principle, since the course of inquiry is unpredictable and since our modal judgments about the capacities of science are not very reliable (see Ladyman and Ross 2007, 16). But we can say, for instance, that the nature of the forms belongs necessarily to free range metaphysics. Still, it remains the case that metaphysical inquiries that investigate topics necessarily alien to science are free range *insofar as* they proceed independently of science.

The topic of investigation is also important to the extent that it constrains the process of naturalization. Whether there are presently discoverable points of contact between science and some metaphysical theory depends, among other things, on the topic of the metaphysical theory. In particular, it depends on whether the domain of science overlaps with the domain of the metaphysical theory, or whether any scientific evidence is relevant to the metaphysical subject-matter, or whether scientific practices or concepts or heuristics can be usefully applied in our investigations of the metaphysical subject-matter. So the topic of inquiry partly determines whether some metaphysical theory *can* be naturalized or not. Still, whether it *is* naturalized or not depends on the actual methods used in its formulation.

<sup>40</sup> Nothing hinges on my specific choice of examples. I select these examples only because they come from prominent metaphysicians; any number of other examples would be equally illustrative.

principles! As I will argue below, free range constraints can be perfectly well-motivated, and to say that they are free range is certainly not to say they should be abandoned.

The second example is Judith Jarvis Thomson's (1998) mereological discussion of the statue and the clay. Thomson proceeds by considering a number of examples, making intuitive observations about them, and accepting or rejecting mereological theses on the basis of those intuitions. For instance, she reports: "Some philosophers... conclude that artifacts cannot undergo replacement of any part, and others that there are no artifacts at all. These [views] strike me as *weird*" (original emphasis, 1998, 153). According to Thomson, a more natural view takes the following intuitive considerations on board:

If you get a new windshield wiper for your car, then in one way, of course, your car is not the same: it has a windshield wiper it formerly did not have... We might say that the car isn't *the same*, for it has changed — but it is *it*, the same car, that has changed. I think we had better agree, and thus that we had better reject the Identity Thesis... it seems as plain as day that we do ordinarily think and say that artifacts can and often do undergo replacement of parts. (1998, 152-3)

She takes these everyday intuitive judgments to be an important source of evidence for or against certain metaphysical theses. She also claims that in general, "Philosophy should not depart more than it absolutely has to from what we ordinarily think and say" (1998, 153). That is, our argumentative conclusions and broader theories should respect our everyday intuitive judgments. So Thomson's primary aim seems to be to preserve ordinary intuitions. Her thesis likely also receives support from its explanatory role. To the extent that Thomson's claims are constrained primarily by intuitive plausibility and explanatory power, and not by science, her (1998) discussion of the statue and the clay is an example of free range metaphysics.

Lastly, Kit Fine's (2001) treatment of grounding is also a clear example of free range metaphysics. Fine's epistemology of grounding is as follows:

[There are] two main sources of evidence for making judgements of ground. The first is intuitive. We appear to be in possession of a wealth of intuitions concerning what does or does not ground what... The other main source of evidence is explanatory in character... a system of grounds may be appraised, in much the same way as any other explanatory scheme, on the basis of such considerations as simplicity, breadth, coherence, or non-circularity. Perhaps the most important virtue in this regard is explanatory strength... (2001, 21-2)

The evidential sources that Fine lists above are logical (coherence, non-circularity), aesthetic (simplicity), and psychological (intuitiveness, explanatory strength). Insofar as Fine's method places only these logical, aesthetic, and psychological constraints on theoretical content, his (2001) metaphysics of grounding is an example of free range metaphysics. So much for examples.

### **I.3 Failure to Meet Epistemic Criteria**

I will argue that free range metaphysics is epistemically inadequate. It is epistemically inadequate because it fails to meet the epistemic principles that I established in Chapter 2. First, the constraints on its content are not sufficiently robust.<sup>41</sup> That is *not* to say we should reject or eschew all of those constraints. Some of them may be well-motivated or useful for the purposes of

---

<sup>41</sup> Karen Bennett shows that in some metaphysical debates "there are *few* grounds for choosing" between rival theories, because measures of simplicity trade off against one another and because the problems that arise for one arise for the other in one form or another (2009, 73). In those cases, she claims, we have "little justification" for believing either view (2009, 42). Her claim is localized to the particular metaphysical debates she takes pains to describe. My claim here is broader. Moreover, my argument focuses not on trade-offs of simplicity or multiply-applicable problems, but on the *weakness* of the theoretical constraints on free range metaphysics.

theory selection (though Bennett 2009 argues that theory-selection criteria don't really help in certain cases). Rather, it is to say that the constraints should not stand alone. Second, satisfying these constraints does not secure epistemic warrant, because the constraints permit into the theory both demonstrable falsehoods and claims that we have little reason to accept.<sup>42</sup> I will demonstrate both of these points by examining the constraints in turn.

### I.3.1 Consistency

The consistency constraint on theoretical content requires that none of a theory's theses contradict one another. While the examples above do not explicitly involve a consistency constraint, the constraint is usually implicit. We almost always require theoretical content to respect a consistency constraint. The consistency constraint is highly permissive. That is because, for any proposition  $p$ , the set of propositions consistent with it is infinitely large. The demonstration is simple — just add disjuncts to  $p$ :  $p$  or  $q$ ,  $p$  or  $r$ ,  $p$  or  $s$ , and so on to infinity. One could equally well add conjuncts, just in case the conjunct neither *is* nor *implies*  $\neg p$ . Further, the consistency requirement does not at all constrain which theses we *initially* countenance. It offers no guidance regarding what goes into the theory in the first place; once we have accepted some theoretical claims, it rules out only the theses that strictly contradict them. So the consistency constraint fails to be robust.

Moreover, satisfying the constraint also fails to secure epistemic warrant, because it permits demonstrable falsehoods and theses we have no reason to accept. For instance, I might

---

<sup>42</sup> Kriegel (2013) argues, similarly, that in revisionary metaphysics, what we take to be theoretical virtues (including, *inter alia*, simplicity and intuitiveness) fail to be truth-conducive. However, my present concern is not directly with truth-conduciveness, but with the aptness of such virtues (or constraints, in my terminology) to constrain theoretical content and secure epistemic warrant.

have a theory according to which the world contains a single apple and nothing else. This theory is internally consistent, since it contains just one self-consistent thesis. But the theory is demonstrably false — look around! — and I have no reason to accept it.<sup>43</sup> Such examples come easily and they show how weakly the demand for consistency constrains theorizing.<sup>44</sup> Again, this does not mean that the principles aren't warranted in their own right or that they should be abandoned. It is simply to point out a feature of them that I have shown (in Chapter 2) to be relevant to the epistemic adequacy of theories.

### I.3.2 Simplicity

Aesthetic norms favouring simplicity are *prima facie* a good deal more robust than the consistency constraint. While these norms are not as ubiquitous as the consistency constraint, they are also frequently assumed. Where they operate — in Fine's method, for instance — they palpably constrain theorizing. In particular, Ockham's Razor precludes theoretical claims that posit entities beyond necessity. However, as William Wimsatt points out, the necessity clause of Ockham's Razor is open to interpretation, so much so that "With the right standards, one could remain an Ockhamite while recognizing a world which has the rich multi-layered and interdependent ontology of the tropical rain forest" (2003, 1). For instance, one could construe the necessity clause in terms of *explanatory* necessity. On that construal, a posit would be necessary if it filled some

---

<sup>43</sup> Of course, the success of the demonstration hinges on our accepting the evidence of our senses. I assume that we should. To do so is not to beg the question, since I am not concerned with whether we should accept empirical evidence in the formulation of metaphysical theories. Of course we should. That is to say, I don't countenance skepticism here.

<sup>44</sup> And this shouldn't be surprising! Certainly no one would have thought that consistency alone secures truth (I thank David Papineau for pointing this out). I grant this. At this juncture, I just wish to highlight how individually weak the free range constraints are – I will later consider whether they are jointly robust.

explanatory gap. And an abundant number of theses, positing an abundant roster of entities, might provide us with the resources to meet that explanatory need — say, by allowing us to explain reference (as the Neo-Fregeans) or intentionality (as Meinong). Now, there are lots of metrics of simplicity (see Kriegel 2013). And, as Bennett points out, gains on one metric of simplicity can result in losses on other metrics, so that “we are just riding a see-saw—fewer objects, more properties; more objects, fewer properties” (2009, 65). All this goes to show that demands for simplicity leave a good deal of wiggle room. While the Ockhamite says that one should have good reasons for positing things, there can be all sorts of good reasons, and some reasons permit all sorts of posits and accompanying theses. So Ockham’s Razor and other norms favouring simplicity fail to be robust on certain interpretations.

Regardless of how one interprets such norms, their satisfaction still fails to secure epistemic warrant. The theory I mentioned above, according to which the world is populated only by a single apple, is simple in the sense that its ontology is sparse. But as I noted above, it is demonstrably false and I have no reason to accept it. Note also that the theory is *both* consistent *and* simple, so adding the two constraints together did not remedy the lack of robustness or epistemic warrant.

In Chapter 2, I argued that considerations of sparsity give us a reason to accept one of the constraint principles — namely the weak constraint principle, according to which the simpler theory is *ceteris paribus* epistemically preferable. How does that square with my argument here that the simplicity constraint secures insufficient epistemic warrant? The idea that simplicity secures some defeasible epistemic preferability does *not* entail that it secures sufficient epistemic warrant. Rather, it only entails that it secures *some* epistemic warrant. I maintain that the warrant it secures is not adequate for the purposes of producing justified theories.

### I.3.3 Intuitive Plausibility

I will construe *intuitions* as immediate, non-inferential judgments (see Devitt 2015).<sup>45</sup> However, I discount immediate perceptual judgments as intuitions. Intuitions in this sense are often used to constrain the content of metaphysical theories.<sup>46</sup> The sorts of intuitions that typically pertain to our metaphysical theories are *philosophical intuitions* — by which I just mean intuitions about matters philosophical, without any commitment to their being a distinct or *sui generis* kind of intuition. More specifically, they are *metaphysical intuitions*, which I take to be a species of philosophical intuition about the world and its workings.

In the context of free range metaphysics, the intuition constraint is more robust than the previous constraints, because metaphysical intuitions suggest a fairly particular conception of the world and its workings, so that one dismisses as strange a good many theses. Moreover, with the intuition constraint comes a requirement for adequacy to common sense and everyday observational facts. As a result, it rules out crank theories like the apple theory. However, it is not clear how much guidance my metaphysical intuitions really offer: while it seems to me that the world contains more than a single apple, my intuitions about, say, Swampman are less clear-cut. Moreover, intuitive plausibility may not really be so robust a constraint, considering that philosophical intuitions more broadly can be pushed around in various ways. Some empirical evidence suggests that philosophical intuitions vary in accordance with: our theories (Machery and Stich 2012); our cultural background, socioeconomic status, and educational background

---

<sup>45</sup> There is, of course, a whole literature on intuition (see Chudnoff 2014 for an overview). I don't mean to commit to any view of what intuitions *are* or what the term *means* — what I give here is just a stipulative definition.

<sup>46</sup> As we have seen, Dorr (2010) claims that intuition-talk by metaphysicians is a kind of humble rhetoric that signals assumptions, rather than being evidential. See Eklund (2013) and Maclaurin and Dyke (2012) for replies.



(Machery, Mallon, Nichols and Stich 2004; Nichols, Stich, and Weinberg 2003; Weinberg, Nichols & Stich 2001); and with other non-truth-tracking factors, like the wording and context of survey questions (Andow 2016; Sinnott-Armstrong 2008; Petrinovich and O’Neill 1996), the order of presentation of cases (Schwitzgebel and Cushman 2012; Swain, Alexander, and Weinberg 2008; Weinberg, Gonnerman, Buckner, and Alexander 2010; Wiegmann, Okan, and Nagel 2012), and emotional affect (Nichols and Knobe 2007), including disgust (Wheatley and Haidt 2005).<sup>47</sup>

But more pressing than the question of *how much* guidance my philosophical intuitions offer, is the question of *why* I should take them as a guide in the first place. When considering the epistemic status of philosophical intuitions, we should consider their source. But since philosophical intuitions are heterogeneous and include intuitions about moral, modal, logical, and linguistic matters, among others, there are potentially many sources. Alvin Goldman (2007) suggests that philosophical intuitions have their source in personal concepts. Timothy Williamson (2007) claims that the source of *modal* judgment is a background of experience and an accompanying folk physics, together with a capacity for ordinary counterfactual reasoning. Some of the experimental results surveyed above suggest that some philosophical intuitions are partly the product of cultural learning. Shaun Nichols and Josh Knobe (2007) also suggest a number of possible psychological sources of moral intuition, including “the distorting effects of emotion and motivation”, an “underlying system for making responsibility judgments”, and “an encapsulated module” (2007, 678-679). Similar models could explain metaphysical intuition in particular, though metaphysical intuitions are probably less vulnerable to emotional affect. So, at the minimum, metaphysical intuitions could have some or all of the following sources: concepts,

---

<sup>47</sup> See Devitt (2012) for a critical response to some of this work.

experience, counterfactual reason, cultural learning (including folk theory), emotion and motivation, and an underlying system or module for making metaphysical judgments.

From an evolutionary perspective, it is doubtful that we have an underlying system or module responsible for truth-tracking metaphysical intuitions about the deep underlying structure of reality, since such intuitions would have been neither necessary nor beneficial for the purposes of survival and reproduction (see Ladyman and Ross 2007, 2).<sup>48</sup> As for the other potential sources, it is not clear why *any* of them should be considered evidential with regard to the deep structural truths of reality. For the most part, personal concepts, cultural learning, emotion, and personal motivation are not relevant to metaphysical truth. These sociological and psychological forms of evidence fail to speak to most metaphysical questions, *qua* questions about *the world*. As for ordinary experience and counterfactual reason, while it is in general a reliable source of evidence, it does not speak to many deep metaphysical questions. My ordinary experience is of macroscopic objects like tables and cats; it tells me very little, if anything, about the underlying nature of the world. In sum: metaphysical intuitions have a variety of potential psychological sources, but none of the ones I have considered would make them good sources of evidence about traditional metaphysical questions.

Moreover, as Ladyman and Ross (2007) point out, metaphysical intuitions have a poor track record when it comes to acquainting us with metaphysical truths. While our metaphysical

---

<sup>48</sup> One might object: it is not clear that there is a dedicated module for scientific intuition, either — a module that would have been beneficial for individual survival and reproduction. I thank an anonymous reviewer for the objection. I agree that there is not obviously a *sui generis* module for scientific intuition. But scientific theories do not rest on some scientific analog of metaphysical intuition. While some metaphysical speculation rests on metaphysical intuition, scientific theory typically does not bottom out in scientific intuition. Rather, it rests on evidence that we gather using capacities that we *do* have evolved neural systems for: perceptual systems, systems that enable action and intervention, and so forth.

intuitions are perfectly adequate for the purposes of our everyday activities, science tends not to vindicate them. Ladyman and Ross point out, “science, especially physics, has shown us that the universe is very strange to our inherited conception of what it is like” (2007, 10). For instance, prior to various scientific and mathematical developments, “metaphysicians confidently pronounced that non-Euclidean geometry is impossible as a model of physical space, that it is impossible that there not be deterministic causation, that non-absolute time is impossible, and so on” (2007, 16). Ladyman and Ross don’t specify which metaphysicians they have in mind. Kant’s defence of the necessity of Euclidean geometry springs immediately to mine. But the curvature of spacetime, the possibility of indeterministic causation, and the framework-relativity of the present moment were all scientific *revelations* — they overturned aspects of our inherited view of the world. Suppose that those aspects rested primarily on metaphysical intuition. This is debatable, but let’s suppose it for the moment. If so, then Ladyman and Ross’ argument does not cut against what I have called free range metaphysics in particular, but against metaphysical intuition more broadly, as well as the forms of inquiry that depend on it (including free range metaphysics). The point is, then, that the world has proven to be counterintuitive. If so, then metaphysical intuitions should not be considered evidential.<sup>49</sup>

Still supposing that the scientific revelations that Ladyman and Ross cite overturned aspects of our worldview supported primarily by metaphysical intuition, one might object: that just shows the fallibility of our metaphysical intuitions.<sup>50</sup> But science and scientific evidence are

---

<sup>49</sup> For further arguments against the evidential value of metaphysical intuitions, see Kriegel (2013).

<sup>50</sup> For this point, I thank an audience member at my talk “Keep the Chickens Cooped: The Methodological Inadequacy of ‘Free Range Metaphysics’” at the *Canadian Philosophical Association Annual Congress* in Ottawa, May 31, 2015.

similarly fallible, as the history of science shows.<sup>51</sup> Moreover, we shouldn't expect sources of evidence to be infallible. I grant the latter two points. However, I wish to stress that science falsifies metaphysical intuitions not just occasionally, but *regularly*. Here we may discharge our above supposition — even if the examples above don't all cut against metaphysical intuition, there are plenty more that do. Ladyman and Ross, citing Lewis Wolpert, give a number of further examples, including that “[m]ost people... are astounded to be told that there are more molecules in a glass of water than there are glasses of water in the oceans” (2007, 11). Andrew Shtulman and Kelsey Harrington (2016) identify a number of common sense metaphysical intuitions that must be unlearned during the process of science learning, including that “heat is... an immaterial substance that flows in and out of objects”, “evolution is... a process that guarantees organisms the traits they need in order to survive”, and “objects... move only if imparted an internal force, or impetus, and will remain in motion until that impetus dissipates” (2016, 119). Shtulman and Harrington also mention some scientific theses that overturn our metaphysical intuitions, including that air is composed of matter and that humans have evolved from sea-faring creatures (2016, 121). There are abundantly many examples of science being at odds with metaphysical intuitions. We might even think this feature is built into science, since the extent to which science correctly makes *surprising* predictions is taken to be a measure of its success. The point is that, with respect to metaphysical truth, metaphysical intuitions continually fail us, so we should think they are systematically unreliable. From their systematic unreliability, we should draw the following conclusion and take its methodological implications to heart: we are just not very good at intuiting the nature of the world. Metaphysical intuitions simply do not earn their keep. To the extent that

---

<sup>51</sup> I thank an anonymous reviewer for raising this point.

we should not trust unreliable sources of evidence, and absent some grounds for thinking that we're improving in this respect, we should assign metaphysical intuitions no great evidential role.<sup>52</sup>

### **I.3.4 Explanatory Power**

Some metaphysicians, like Fine (2001), appeal to explanatory power as a criterion for the evaluation of metaphysical theses. In the sciences, explanatory power can be an important indicator of the credibility of a hypothesis (for examples, see Kitcher 1981). And, as Ladyman and Ross point out, “[s]ome metaphysicians have realized that they can imitate science by treating their kind of inquiry as the search for explanations too, albeit in a different domain” (2007, 17). Many metaphysical theses are produced by abduction and so are meant to explain certain evidence. For instance, the claim that there are non-existent objects is meant to explain how we can have intentional attitudes toward Sherlock Holmes and square circles (see Crane 2013 and Priest 2014). In the following section, I will argue that in the context of free range metaphysics, the explanatory power constraint is both permissive and insufficient to secure epistemic warrant. Recall, however, that I don't take that conclusion to show that we should abandon the constraint.

*Prima facie*, very many things can putatively explain a set of evidence.<sup>53</sup> For instance, the evil demon hypothesis putatively explains my phenomenal states. The hypothesis suggests that for each of my phenomenal states, I have it because the demon wills it so. This putative explanation

---

<sup>52</sup> Note that this argument is not probabilistic. It makes no claim about the probability of current or future metaphysical intuitions turning out to be false. It simply denounces a form of evidence on the basis of its systematic unreliability. So the base rate fallacy does not threaten here.

<sup>53</sup> I say ‘putatively explain’ in case the reader thinks explanations must be factive. If they must be factive, then there is an epistemic problem with regard to the conditions under which we can *know* something to be an explanation. But at any rate, what we care about is *putative explanations*, since it will be *those* that constrain our theories. We will permit into the theory what we *take to be* explanatory.

has some virtue: it unifies a host of myriad phenomena under a single, simple explanation. Since I could dream up any number of other explanations like it, the requirement that something be explanatory doesn't seem particularly robust.

But the evil demon hypothesis may fail to be a *good* or *powerful* explanation. So we have to consider what makes a metaphysical explanation good or powerful. These are loaded questions, too grand to settle here, for which there may not be a single answer (Colombo 2016). So let me remain neutral on the precise account of explanation and instead consider some well-known candidate measures of explanatory power.

Peter Lipton (2004) suggests that explanatory power can be gauged in terms of i) the extent to which the explanation renders the explanandum intelligible, or ii) how likely the explanation is given the evidence. He calls these, respectively, the *loveliness* and *likeliness* of the explanation (2004, 59). Let's consider *loveliness* first. According to Lipton, the loveliest explanations "provide the most understanding" (2004, 59). But this needs spelling out. How does an explanation facilitate intelligibility or understanding? Among other things, it could be a matter of providing a relevant causal story (Salmon 1984) or of unifying knowledge (Kitcher 1981, 1989). While Salmon and Kitcher account for *scientific* explanation, the central thrust of their views might help us flesh out how metaphysical explanation can be lovely in Lipton's sense. In the sections below, I will consider both options. I will find that on both the causal and the unificationist accounts, the explanatory power constraint operates against a background of empirical belief. The free range metaphysician may argue that the empirical background helps to robustly constrain and support her theories. However, I will show why that is not so.

Suppose first that explaining is a matter of providing a relevant causal story. For any given explanandum, I can tell all sorts of relevant causal stories. For instance, for each of my phenomenal

contents, I can explain that content by saying the evil demon caused it, or that God did, or gremlins, or whatever. Any number of relevant putative causes could explain my phenomenal experience. So for a causal story to give us a *good* or *powerful* explanation, it should be *plausible given what we know*. Given what we know, the evil demon hypothesis and its kin are not plausible causal stories. They are certainly logically possible and *consistent* with the available empirical evidence, but our evidence does not support them. We have a good deal of knowledge about the world, and our knowledge of the world provides us with little, if any, evidence from reliable sources of supernatural creatures such as evil demons. Of course, this line doesn't defeat the skeptic, but it is a common sense response that one might find attractive. At any rate, notice that the *given what we know* clause brings into the picture a background of empirical belief. So on this interpretation, our loveliness criterion is empirical. And the empirically-informed background of belief may be said to constrain the causal stories we countenance.

And it's not just this causal interpretation of explanatory power that builds in a background of empirical belief. Suppose it's *the unification of knowledge* that facilitates understanding and thereby makes an explanation lovely. If so, an explanation facilitates understanding by fitting into our systematic picture of the world. Now, in Kitcher's unificationist view, the 'systematic picture' is the set of sentences currently accepted by a scientific community (1981, 512 & 519). If we follow Kitcher on this matter, then the explanatory power criterion becomes a scientific constraint, not a free range one (recall that the constraints on free range metaphysics by definition exclude scientific constraints). Since we are currently evaluating the tenability of explanatory power as a constraint on free range metaphysics, Kitcher's particular unificationist view is not of present concern. However, one could have a unificationist theory of explanation that does not build in science. One could hold a view according to which something is explanatory relative to a set of

folk beliefs held by an individual or shared by a community. To the extent that our folk picture of the world is based largely on our empirical observations, this way of cashing out the explanatory power constraint would make it inherently empirical. So plausible renderings of both the causal and the unificationist understandings of explanatory loveliness makes them empirical.

The same goes for likeliness. My judgment that an explanation is *likely* could be one of two things: either 1) an intuition or 2) the output of a statistical inference. Either way, empirical information will constrain my judgment. If my judgment is intuitive, my intuition is formed against a background of empirical belief. If my judgment is the output of a statistical inference, then an empirical background of belief helps to determine the priors that will, ideally, figure into my reasoning. So if explanatory power is a function of likeliness, then the explanatory power constraint is empirical. The upshot of the discussion so far is this: however we understand explanatory power, the explanatory power constraint is empirical.

Note that this implies that free range metaphysics, insofar as it is beholden to an explanatory power constraint, is not *a priori*. So, while it may be tempting to use ‘a priori metaphysics’ as a quick gloss for the sort of metaphysics I am criticizing, that is just not the right joint to carve at. Setting free range metaphysics aside for a moment, we might wonder whether the above paragraph suggests that there is no *a priori* metaphysics whatsoever. It does imply that any form of metaphysics that requires its contents to be explanatorily powerful is not *a priori*. While it is certainly possible for metaphysics to proceed without any explanatory aims — imagine some metaphysical theory about an esoteric topic removed from experience that we just aim to make an internally coherent system — I will remain neutral on whether it ever does so proceed. It would be an interesting research program to examine the explanatory demands on current metaphysical theories.



The free range metaphysician may feel a glimmer of hope at this juncture. She may feel that, having established that the explanatory power constraint is empirical, she has a foothold to defend her method. Surely, she thinks, the background of empirical belief makes the explanatory power constraint robust, since it limits the number of acceptable metaphysical explanations. While the empirical background cannot decisively rule out the truth of evil demon hypotheses (no evidence does), it does make them implausible and therefore precludes her from countenancing them as good or powerful explanations. The world with which our hopeful metaphysician is experientially acquainted seems not to include things like supernatural demons or gods or gremlins. So the background of empirical belief will preclude her from finding the crank hypotheses she dreams up explanatorily powerful. Moreover, the hopeful metaphysician also thinks that she can secure epistemic warrant on the basis of the empirical credentials of the explanations she finds powerful. Surely, she thinks, free range metaphysical explanations receive warrant by being held accountable to the tribunal of experience. So the hopeful metaphysician thinks that the explanatory power constraint, *qua* empirical constraint, can robustly constrain her free range metaphysical theories and secure support for them.

My reply is two-fold. First, in the context of free range metaphysics, the explanatory power constraint operates against a *folk-theoretic* background of empirical belief. That is, it operates against a background of belief that is only minimally informed by science. Outside free range metaphysics, explanations are frequently judged to be powerful against a fuller background of scientific knowledge. But once the institutional outputs of science are playing any more than a nominal constraining role, by definition we no longer have free range metaphysics.<sup>54</sup> So we have

---

<sup>54</sup> That is why scientific theories, such as Darwin's evolutionary theory, whose success is primarily a function of their explanatory power, are not free range theories.

at best a kind of folk physics in the background. But folk theory shouldn't constrain metaphysics, for reasons we have already seen. Namely, the world tends to frustrate our intuitive or folk-theoretic conception of it. So *qua* free range theoretical constraint, explanatory power fails to secure epistemic warrant, notwithstanding the fact that it operates against a partly empirical background of belief.

Second, even operating against an empirically-informed background, the explanatory power constraint *still fails to be robust and to secure epistemic warrant*. Let's dispense with the unificationist understanding of explanatory power first, since the explanation of its failure differs from the others. First, for the purposes of metaphysical inquiry, we should not be interested in unifying folk theory, since it is partly based on intuitions. Second, all sorts of metaphysical explanations, including crank explanations that I have no reason to accept, can unify a system. I can derive all the constituents of my system of belief from a God hypothesis, or an evil demon hypothesis, and so on. So on a unificationist construal, the explanatory power constraint fails to be robust and to secure epistemic warrant.

On the other construals of explanatory power, judgments of explanatory power collapse into intuitive judgements, which we have seen to be both permissive and epistemically dubious. Which *causal stories* we find plausible given our background beliefs is an intuitive matter. Many people find it plausible, even against an empirically-informed background of beliefs, that angels, devils, and gods figure into explanatory causal stories. Likewise, which explanations we find *likely* can be a straightforwardly intuitive matter — a matter of which ones we find plausible given our evidence. Even judgments of the statistical probability of putative metaphysical explanations rely on intuition. Although we assign priors against a background of empirical knowledge, to the extent that the background sometimes offers little guidance, we set those priors on the basis of intuition.

So judgements of explanatory power, understood in terms of relevant causal stories or in terms of likeliness, depend crucially on intuitive judgment. And as we have seen, intuitive judgment fails to be robust and to secure epistemic warrant. My point is this: no matter how we cash out the explanatory power constraint on free range metaphysics, *notwithstanding* its empirical status, the constraint fails to be robust and to secure epistemic warrant

### **I.3.5 No Joint Sufficiency**

So far I have shown only that the constraints traditionally taken to constrain metaphysical theorizing are not individually adequate constraints. Might they be jointly sufficient? Suppose a theory is logically consistent, simple, intuitive, and explanatorily powerful. Is that enough for robust constraint or epistemic warrant? I have demonstrated that each theoretical constraint is extremely weak. Adding up these extremely weak constraints would amount to still minimal theoretical constraint. We have already seen that adding up some of them did not help matters, since the apple theory was both simple and consistent. Take those two constraints and add to them the demands for intuitive plausibility and explanatory power. We have seen that on several conceptions of explanatory power, judgments of it are intuitive judgments. Where that is the case, the explanatory power constraint collapses largely into the intuitive plausibility constraint. That is, most of the constraining work is done by intuitions. So we now have essentially three constraints: simplicity, consistency, and intuitive plausibility. Together, the first two are highly permissive. As we have seen, there are reasons not to place the third constraint on our metaphysics at all. But setting that aside, we have also seen that the intuitive plausibility constraint is permissive to the extent that our intuitions can be pushed around in various ways. So imagine the range of theories that could, actually or potentially (with some manipulation of our intuitions), satisfy our

requirements for simplicity, consistency, and intuitive plausibility. It's enormous. This indicates that the constraints on free range metaphysics are not jointly robust. And in some ways, the state of metaphysics reflects this. Though disagreement is a feature of most epistemic life — in the sciences, too — the deep and persistent disagreements that occur in free range metaphysics, and the sheer number of theoretical alternatives on offer for each subject matter, are symptoms of this insufficient constraint. So the constraints on free range metaphysics, both individually and jointly, fail to be sufficiently robust or to secure epistemic warrant. Therefore free range metaphysics is epistemically inadequate.

#### **I.4 Collateral Benefits**

Before drawing a normative conclusion from the various epistemic deficiencies I have pointed out, I should pause. As Steve French and Kerry McKenzie point out, “whatever exactly the problem with contemporary metaphysics is taken to be... the appropriate *reaction* to it... has to be considered carefully” (2012, 44). That is because, notwithstanding its deficiencies, free range metaphysics may have some value that we should recognize and preserve. I will now consider some candidate sources of value — that is, some possible collateral benefits of free range metaphysics.<sup>55</sup>

##### **I.4.1 Clarity**

In his (2010) review of *Every Thing Must Go*, Cian Dorr comments that Ladyman and Ross “have missed what is best and most distinctive about the tradition they set themselves against: its

---

<sup>55</sup> Note that my use of the phrase ‘collateral benefits’ here echoes Maclaurin and Dyke’s use in their (2012).

gradual raising of the standards of clarity and explicitness in the statement of metaphysical claims” (2010, np). Just as analytic philosophy more broadly, with its attention to logic and language, increased the standards of clarity in the statement of philosophical claims, analytic metaphysics has increased the standard of clarity that we hold metaphysical claims to. According to Dorr, “[m]uch of what is distinctive about the analytic way of doing metaphysics is meant to guard against the danger that we might accidentally lapse into nonsense, or launch into disputes that turn out to be merely verbal” (2010, np). To that end, analytic metaphysicians focus on theses “that can be stated using familiar everyday words” (2010, np). When they introduce a technical term, they “propose logical constraints on the new vocabulary, and attempt to draw connections between it and questions expressed in more familiar terms, in the hopes of thereby imposing enough discipline on its use to fend off the charge of unintelligibility” (2010, np). For instance, to their great benefit, metaphysicians have required that the terms *fundamental* and *derivative* be accompanied by an “adequately expressive language” for talking about more fundamental things without mentioning derivative things (2010, np). As a result of their “patience for fine distinctions and quibbling objections”, analytic metaphysicians have developed “techniques... for stating claims clear and explicit enough to be worthy targets of argument... [and] crafting a view coherent and explicit enough for arguments to get any grip” (2010, np). The claim is not just that free range metaphysics *is really clear*. To that, I could respond that clarity for its own sake has limited value. Further, the claim is not that clarity does not, or cannot, occur in other domains. That would be patently false. The claim is, rather, that the discipline of analytic metaphysics has produced something of value: techniques for the formulation of claims that meet a comparatively high standard of clarity. One might think that having done so somewhat mitigates against the epistemic failings that I have detailed.

### I.4.2 Fruitfulness in Science

Free range metaphysics is sometimes fertile ground for science. Philosophy more broadly sometimes plays what Peter Godfrey-Smith (2013, 2014) calls an *incubator* role with respect to the sciences. He explains, “In its relation to science, philosophy has often also functioned as an ‘incubator’ of theoretical ideas, a place where they can be developed in a speculative way while they are in a form that cannot be tested empirically” (2014, 1). He gives a number of examples: associationism; much of the current theoretical framework in cognitive psychology and linguistics; the embodied approach to cognition; the Bayes net approach to causal relations; the modern development of computers; and even the more scientific aspect of Marxism, the origin of which is to be found in Hegel (2013, 2-3).<sup>56</sup> One might think that free range metaphysics shares this incubator role. On such a view, metaphysical theories are speculative sketches that are not sufficiently worked out to be empirically tractable; nonetheless, they sometimes give rise to questions of scientific interest and tractability. It is not unlike the way that theoretical physics sometimes incubates experimental physics — as it did in the case of general relativity. While philosophy “has no monopoly” on this incubator role, and nor does free range metaphysics, “[i]ncubation of new ideas is undeniably important” (2013, 3). One might argue that whatever normative conclusion I draw, it should preserve this incubator function.

---

<sup>56</sup> One might disagree that Bayes nets count as a clean example of philosophy incubating science, since variants of Bayes nets were used in statistics, genetics, and econometrics well before philosophers carved out the field of study. I thank David Papineau for pointing this out.

### I.4.3 Fruitfulness in Scientifically Engaged Philosophy

Steve French and Kerry McKenzie (2012) articulate a further collateral benefit of free range, or, as they call it, *fantasy* metaphysics (2012, 43). Although they share Ladyman and Ross' (2007) dubiousness of such metaphysics, they resist their conclusion that it should be discontinued. In fact, they argue that proscribing such metaphysics would be “counterproductive” (2012, 44). That is because, according to French and McKenzie, free range metaphysics is enormously beneficial to those who engage philosophically with science. It is a toolbox; it provides conceptual and formal frameworks useful in the formulation of scientifically engaged philosophical theories. As they put it, “the products of analytic metaphysics can be regarded as available for plundering... in order that we might exploit them for our own purposes” and avoid having to “develop the appropriate resources from scratch” (2012, 44). More recently, French calls this the *Viking Approach* to metaphysics, in which metaphysicians play the role of “hapless peasants, happily tilling their fields of compositionality and ontological dependence, before being pillaged by ruthless realist marauders” (2014, 50). For instance, ontic structural realists pillage or plunder metaphysics when they invoke notions of “‘modal structure’... fundamentality and the metaphysics of relations” (2012, 44). Further:

The growing literature on ontological dependence... is proving useful in expressing the core metaphysical claim of ontic structuralism, namely, that physical objects are ontologically secondary to structures... A form of truthmaker theory might also be deployed in order to articulate the eliminativism about objects that ‘radical’ ontic structural realism endorses... Even the work of Lewis – a philosopher who is often pilloried for his lack of engagement with science – has been summoned in defense against the triviality objection to structuralism. (French and McKenzie 2012, 44)

So free range metaphysics furnishes us with conceptual and formal tools useful in the articulation of scientifically responsible views. French and McKenzie draw an analogy to the usefulness of pure mathematics in physics. Just as non-Euclidean geometry was useful to Einstein,

so it is useful to eliminative structuralists that there [is an available] theory of dependence compatible with the elimination of the dependent entity. And just as it was useful to the development of particle physics that the theory of Lie groups was largely completed by the time the appropriately high-energy regimes could be probed, so it is beneficial to the defender of the Everett interpretation that [there is] a theory of personal identity that makes decision-making make sense in branching universes... (2012, 44)

Just as pure mathematics can turn out to be useful in physics, free range metaphysics can turn out to be useful in scientifically engaged philosophy. Because free range metaphysics can offer tools useful to our more legitimate theoretical enterprises, “we, as naturalistic philosophers, have to *value* scientifically disinterested metaphysics” — or at least recognize its potential value (2012, 45). To the extent that free range metaphysics produces useful tools, it is “very difficult to oppose scientifically disinterested metaphysics *tout court*” (2012, 45). And, French and McKenzie note, since the course of inquiry is unpredictable, “It seems folly to try to predict in advance what will or will not prove useful to us in the course of time” (2012, 45). Since free range metaphysics cannot produce useful tools in a purposive, premeditated, directed manner, it must do so by happy accident. And for that reason, one might say, with French, “Run free, metaphysicians!”<sup>57</sup> That is,

---

<sup>57</sup> French made this remark following his talk “Between Humeanism and Dispositionalism; or, How to Construct a Modal Framework for Modern Science by Appropriating Metaphysical Devices” at the conference *New Trends in the Metaphysics of Science* on December 17, 2015.



keep free range metaphysics *free range*; allow it to produce a thousand flowers, with the hope that some of them turn out useful.<sup>58</sup>

## **I.5 Normative Conclusion**

The normative conclusion that I draw here has a good deal to keep in balance. On the one hand, it should reflect and remedy the epistemic deficiencies I have pointed out. But on the other hand, if we take free range metaphysics to have collateral benefits that are sufficiently valuable, it should also respect and preserve them.

Just how valuable are these collateral benefits? Well, free range metaphysics is not *unique* in having them. Other activities — including analytic philosophy broadly construed, theoretical physics, and pure mathematics — yield techniques conducive to great clarity, incubate scientific activity, and produce useful formal tools, respectively. Now, one might argue that free range metaphysics accomplishes one or many of those tasks uniquely well. But I will not consider that argument here. My interest is just in the fact that it has benefits at all. While other intellectual activities may produce benefits of the same type, free range metaphysics produces a number of token benefits, and has the potential to produce more. To the extent that these token benefits aid human inquiry, free range metaphysics has value.

Since consigning inquiries to the flames is a grave matter, we should do so only when the inquiry does much more harm than good. So let us take stock of the harm free range metaphysics does. As French and McKenzie point out, metaphysicians believe themselves to accomplish a good deal more than collaterally aiding human inquiry. They “hold that there is value to their discipline

---

<sup>58</sup> But as Dennett notes, “just remember: if you let a thousand flowers bloom, count on 995 of them to wilt” (2006, 40).

beyond it serving as a production line for constructions that might eventually be used and abused by [other inquirers]” (2012, 45). Metaphysics is “intended to be more than just a toolbox and to itself have some direct bearing on reality” (2012, 43). That is, metaphysics attempts to discern metaphysical truth, or at least to produce justified theories of reality. Bennett helpfully surveys some of the various ways that metaphysicians conceive of their discipline:

They say that metaphysics is the study of ‘being *qua* being’ (Aristotle, *Metaphysics* book IV). They say that it is the attempt to ‘get behind all appearances and describe things as they really are’ (van Inwagen 1998, 11), and that it is the study of ‘what the world is like... as opposed to... how we think and talk about the world’ (Sider 2008, 1 and note 1). They say that metaphysics is ‘inquiry into the most basic and general features of reality and our place in it’ (Kim and Sosa 1999, ix). They say that it is the study of ‘the fundamental structure of reality’ (Sider 2011, 1). And they say that is ‘about what grounds what. It is about the structure of the world. It is about what is fundamental, and what derives from it’ (Schaffer 2009, 379). (2016, 28-29)

Clearly, metaphysicians *standardly* conceive of their discipline as one that acquaints us with truths about the world. But given the epistemic deficiencies I have articulated, metaphysics is not adequately suited to do so. While the outputs of free range metaphysics have some pragmatic value, they lack sufficient epistemic justification. So free range metaphysics is harmful to the extent that its proponents believe it to be an epistemically adequate form of inquiry that produces justified theories about the nature of the world. That is, the harm is in thinking that free range metaphysics has fully satisfactory epistemic credentials over and above its pragmatic ones.<sup>59</sup>

---

<sup>59</sup> Compare van Fraassen (1980), who argues that extra-empirical theoretical virtues are merely pragmatic and therefore don’t provide reasons for belief.

Even humbler metaphysicians who see their discipline as “not about explanatorily ultimate aspects of reality that are actual” but rather, about “*the most explanatorily basic necessities and possibilities*” (Conee and Sider 2005, 203), or as an inquiry into “what there could be” (Lowe 1998, 9), exhibit bad faith. Here I wish to echo Callender (2011), who argues against the value of modal intuitions that are independent of science. He comments, “From the history of science don’t we learn that many ‘impossibilities’ end up possible, and vice versa?” (2011, 41). He continues:

[T]here is no interesting species of metaphysical modality that is largely immune to science. Our modal intuitions are historically conditioned and possibly unreliable and inconsistent. The only way to weed out the good from the bad is to see what results from a comprehensive theory that seriously attempts to model some or all of the actual world. (2011, 44)

And, of course, our best comprehensive theories of the actual world are scientific ones. We shouldn’t take seriously any modal claim generated independently of such a comprehensive theory: “Being connected to a good systemization of the world is either *constitutive* or *symptomatic* of serious possibilities” (2011, 45). Likewise, French and McKenzie argue that “*even if* metaphysics regards itself as the study of the possible, given the central *methodological* role of the actual in systematic modal theorizing and physics’ privileged role within it” metaphysics must be scientifically informed (2012, 56-7). Clearly, free range theorizing is equally dubious in limning the structure of metaphysical modality as it is in describing the structure of the actual world. The metaphysician cannot escape my criticisms by claiming that she deals only in possibilities or in modal structure more broadly. The point stands that the harm of free range metaphysics is its bad faith; any free range metaphysician who thinks her inquiry adequately justifies metaphysical

claims — whether they be claims about the actual world or about modal structure — participates in that harm.

While the harm is serious, some institutional delusion does not so far outweigh the above pragmatic benefits that it warrants consignment to the flames. That is, I cannot in good conscience recommend the discontinuation of free range metaphysics. But I can recommend resolving the discipline's bad faith. The free range metaphysician *need not* be deluded about the epistemic credentials of her discipline. Such delusion is not requisite to achieve the collateral benefits discussed above. Rather than thinking of herself as producing justified theories, the free range metaphysician can think of herself as systematically exploring theoretical space and building coherent systems, in the process of which she sometimes produces beneficial theoretical goods. So the normative conclusion I wish to draw here is this: free range metaphysicians should not deceive themselves by thinking they produce justified theories of the underlying nature of reality.

If we are interested in producing justified metaphysical theories, then we had better keep the chickens cooped — that is, we had better place further constraints on the content of our free range metaphysical theories where possible. It may not always be possible, in which case the free range activity may persist and continue producing its actual and potential collateral benefits. But we should determine where it is possible and constrain accordingly. I will argue in the next section that scientific constraint can be fruitfully added to metaphysics.

## Part II: Scientifically Responsible Metaphysics

### II.1 Definition

The more a metaphysical theory *engages, conscientiously and painstakingly, with the empirical data, theoretical insights, and practices of the current sciences*, the more scientifically responsible it is.

Let me unpack this characterization. First, the language of the definition indicates that scientific responsibility is not a binary. There is no precise threshold above which all theories are equally scientifically responsible, and below which all theories are equally not. Rather, being scientifically responsible is something that admits of degrees — theories can be relatively more or less scientifically responsible than their actual and potential rivals. The more a theory engages with science, the more scientifically responsible it is.

I say that the engagement must be conscientious and painstaking to rule out cases where, as Mark Wilson says, one *mentions* science rather than *using* it (2006, 13) — where one haphazardly peppers one's work with science jargon, but fails to attend to science in its detail and difficulty. I also wish to rule out cases where one simply formulates a metaphysical theory independently of science, checks for consistency as an afterthought, and finds that consistency obtains as a happy accident. Accidental consistency with science does not constitute proper engagement with science, and does not make a theory scientifically responsible.

Metaphysics may engage with science in various ways. *Inter alia*, the modes of engagement include:

- unifying scientific theses (Ladyman and Ross 2007)
- integrating scientific postulates or theses into metaphysical theories in a non-ad-hoc manner

- drawing conclusions of philosophical interest on the basis of scientific evidence
- drawing theoretical conclusions that track the practices of scientists<sup>60</sup>
- interpreting scientific data
- correcting scientists' interpretation of the data where appropriate
- revising metaphysical claims in light of new data or theoretical understanding
- formulating new or revising old metaphysical questions in response to new scientific understanding.

## II.2 Clarifications

In my view, making a domain scientifically responsible is not a matter of co-opting the methods of science. It is not, to borrow Papineau's example, as if philosophers have to start fiddling with particle accelerators (2014, 166). It is primarily a matter of using science as an evidence-base and thereby constraining one's theoretical activity. In other words, scientifically responsible metaphysics is, to borrow a phrase, "philosophically processed' science" (Godfrey-Smith 2012, 99). But I would resist calling scientifically responsible metaphysics *scientific*, insofar as its methods are not distinctively scientific.

Let me also make clear that I do not disavow armchair reasoning. Armchair reasoning has its place both in science and in scientifically responsible metaphysics. In fact, it is, in my view, a primary methodological tool of scientifically responsible metaphysics, and accomplishes many things for it. It reveals the truths of logic and mathematics. More importantly, conceptual work

---

<sup>60</sup> Here I credit Marc Ereshefsky and Ken Waters, who each made the point that naturalistic metaphysics should track scientific *practice* in their presentations at the *Science and Metaphysics: 50 Years of Philosophy of Science* symposium at the *Canadian Philosophical Association Annual Meeting* in Calgary, May 2016.

allows the metaphysician to precisely define her target of investigation. We must have a clear question before we can set about answering it. Further, armchair reasoning allows the metaphysician to think metaphilosophically about her principles of theory construction and criteria of epistemic adequacy, as I do here. Relatedly, given the fact that both scientific and metaphysical theories are underdetermined by their evidence, reason is also important when it comes to theory choice. Armchair considerations may help us choose between metaphysical theories that are on a par in terms of the extent to which they incorporate, explain, and unify scientific evidence. They may also help in the selection of scientific theories, or interpretations of theories, where the theories and their interpretations are equally empirically adequate. When faced with underdetermination, we may appeal to, for instance, coherence, elegance, and other extra-empirical criteria for theory selection.

Furthermore, ampliative inference is instrumental for the construction of metaphysical theories because the substantive deductive implications of science are likely few. Indexing metaphysical theorizing to current science ought to rob the discipline neither of its imagination, nor its autonomy. Metaphysics can do more in relation to science than simply clean up after it. That is, it can do more than merely clarify or unify its theses, though those are among the tasks it may perform. Good metaphysics on my view remains an exploration of theoretical space — it is just that the exploration is tethered by the scientific constraint, and therefore the range of exploration is significantly narrower than it would otherwise be. So the claim is not that armchair reasoning should play no role in metaphysical theorizing, but rather that it is incumbent on metaphysicians who seek justified theories of reality to constrain their speculation by appeal to science.

Now, one might think that I am a verificationist in disguise, because the language of confirmation and disconfirmation is really just a holist-friendly gloss on verification and falsification. However, I reject verificationism, both as a criterion of meaning and as a criterion for the formulation of theses of interest and importance. In my view, the aptness of scientific theses to be confirmed and disconfirmed is important because it makes science a good constraint on metaphysical theories. But scientifically responsible metaphysics does not require that every metaphysical thesis be provable, or posit be motivated, simply by pointing to empirical phenomena or citing scientific evidence. Instead, the view grants some weight to Quinean considerations of holism and coherence, so long as suitably many theses of one's theory receive scientific support.

But here I need an explicit criterion of allowability. That is, I need a rule stating when a thesis is allowable by appeal to holism and when not. Otherwise, anything that *fits* into the theory is allowable, and suddenly our theories are open to God hypotheses and the like.<sup>61</sup> I propose the following criterion: a thesis that lacks direct scientific support is allowable only when it: 1) does not significantly constrain other components of the theory or 2) is itself constrained by components of an overall scientifically responsible metaphysical theory that are themselves supported, ultimately, by science. That is, the thesis has few implications and does not rule out many other theses, or it is supported, ampliatively or deductively, by parts of the theory that have scientific support either directly or indirectly. This criterion mitigates the epistemic risk of countenancing a thesis that does not have direct scientific support, because it either *contains* risk, by preventing it from spreading, or *minimizes* risk, by virtue of chains of evidential support that ultimately trace back to science.

---

<sup>61</sup> I thank Yuval Abrams for this point.



### II.3 Examples

The classic example of a metaphysical topic that has made fruitful contact with science — and a good paradigm case of metaphysics being made scientifically responsible — is the philosophy of time. In particular, the rejection of the notion of absolute time in response to special relativity. Special relativity revealed that what counts as simultaneous is entirely dependent on the movement of the relevant frame of reference, which undermined views, like the A-theory and presentism, that assume a unique present (see Putnam 1967, Sklar 1974, Stein 1968 & 1970, Weingard 1972).

Another example is to be found in discussions of free will that consider quantum indeterminacy. The revelation from quantum mechanics that the causal nexus contains pockets of spontaneity, indeterminacy, randomness, and chaos, falsified a longstanding theoretical opponent to free will: the thesis of causal determinism. Much work in the literature on free will now attempts to discern the implications of quantum indeterminacy for considerations of free will (see Bishop 2002, Hodgson 2002, Kane 1996, Loewer 1996, Margenau 1967). Robert Kane, in particular, argues for a libertarian view according to which agents deliberating between certain actions have greater sensitivity to “micro-indeterminacies at the neuronal level”, which get amplified to the macro-level by “complex chaotic processes” (1996, 130).

A further example is the attempt to draw ontological implications out of quantum mechanics (see Allori 2015; Costa, Lombardi, and Lastiri 2013; Dorato 2015; Esfeld, Lazarovici, Hubert, and Dürr 2014). In particular, some philosophers in the literature argue that quantum mechanical considerations should lead us toward ontic structural realism. That is, they argue that

quantum mechanics is suggestive of an objectless ontology (Esfeld 2004, French 2014, Ladyman and Ross 2007).

Lastly, a good deal of philosophy of science, as well as the burgeoning field of scientific metaphysics, contribute paradigmatically to what I am calling scientifically responsible metaphysics. See for example the biological essentialism and species debates in the philosophy of science.<sup>62</sup> See also Albert (2015), Calosi and Morganti (2016), Dorato and Morganti (2013), Kistler (2010), Ladyman, Ross, and Kincaid (2013), Loewer (2012), Maudlin (2007), Morganti (2008), Ney (2012), Ney and Albert (2013), Norton (2015), and Pradeu and Guay (2015).

#### **II.4 Satisfaction of Epistemic Criteria**

Scientifically responsible metaphysics satisfies my epistemic principles by introducing a scientific constraint on theoretical content. Scientific claims, together with the appropriate auxiliary hypotheses, frequently have observable consequences that we can test against experience. They are frequently confirmed and disconfirmed. I say *frequently* because some scientific claims are physically impossible to test. But in general, science proceeds via the formulation of hypotheses testable under repeated empirical experimentation. When a thesis is disconfirmed, it is expelled from the theory or revised. That's the bouncer in action, throwing people out of the club. Certainly science has disconfirmed a vast number of theses (theses committed to spontaneous generation, caloric fluid, luminiferous aether, and so on), and will disconfirm vastly many more. So by virtue of the disconfirmation of scientific theses, the scientific constraint on metaphysical theorizing is robust. Moreover, when a thesis is confirmed under repeated experimental testing, it

---

<sup>62</sup> See Barker (2010); Devitt (2008); Ereshefsky (1991), (1994), (1998), (2001), (2010a), (2010b); Ereshefsky and Pedroso (2013); Ghiselin (1974); Grant (2011); Hull (1978); Leslie (2013); Mayr (1963), (1970); Okasha (2002); and Sober (1980), (1993).

receives further and further evidence in its favour. So by virtue of the confirmation of scientific theses, the scientific constraint on metaphysical theorizing secures epistemic warrant. By virtue of the confirmation and disconfirmation of scientific theses, a metaphysics that engages with those theses and uses them as evidence is much better constrained and supported than free range metaphysics.

Now, one might argue that the scientific constraint does not rule out crank hypotheses about gods and evil demons and so falls prey to the same criticism I levelled against the constraints on free range metaphysics. Granted, the scientific constraint doesn't rule out crank hypotheses in the sense of demonstrating their falsity. But it does rule them out in the sense that science simply doesn't countenance them. That is, to the extent that god and evil demon hypotheses are not empirically tractable, they are not taken to be scientifically interesting, or the proper subject-matter of science. And this is not merely a sociological fact. Science tends to entertain theses that have novel predictive power — not necessarily in the sense of predicting results of future interventions, but in the sense of predicting formerly unknown evidence. To the extent that god and demon hypotheses do not have predictive power, they fall outside the domain of scientific interest.

That said, not all scientific theses have novel predictive power. For instance, it is not clear that string theory has predictive power (Dawid 2013). Moreover, some theories are taken seriously before their novel predictive power is established (Kitcher 1981). What shall we say about those cases? We might say that science establishes a stock of explanatory patterns legitimized by their place in science as a whole, and that the theses taken seriously by science share features in common with the members of that explanatory stock. So in addition to lacking novel predictive power, crank hypotheses are not countenanced in science because they do not sufficiently resemble the

explanatory patterns that fit into and work in science. Therefore, my conclusion that the scientific constraint is robust holds.

## **II.5 Normative Conclusion**

Because the scientific constraint on metaphysical theorizing satisfies my epistemic principles in being robust and securing epistemic warrant, I conclude that any metaphysics aiming to furnish us with justified claims about the nature of the world should be scientifically responsible.

## **Part III: Objections and Replies**

### **1. Why go beyond science?**

One might wonder why we should bother with scientifically responsible metaphysics at all. After all, it remains a speculative activity that departs substantially from the evidence. In constructing our picture of the world, why not stop at science and go no further?

### **Reply**

If science has a face value, it's a bewildering one. The state of scientific theorizing is a disunified jumble (Ereshefsky 1998, Ruetsche 2011, Waters 2014). So science does not simply hand us a coherent or intelligible picture of the world. But so far as possible, we want a coherent, intelligible picture of the world. We want to more than simply cognize data, theories, and interpretations, but to makes sense of them. As epistemic agents standing in relation to our world and trying to operate within it, we want some worldview that enables understanding and facilitates successful interventions. To that end, we must accept some epistemic risk and go beyond science.

## 2. What role for philosophers?

One might wonder, since scientists are best equipped to understand, interpret, and draw conclusions from scientific results, why philosophers should have a role in scientifically responsible metaphysics. That is, if scientifically responsible metaphysics is an activity in understanding, interpreting, and drawing conclusions from science, why shouldn't we leave the task to scientists? Lowe articulates the objection (which he rejects) as follows:

to the extent that there is a legitimate province of metaphysical inquiry, it is one entirely catered for by the empirical sciences... This being so, there can be no scope for a distinctly 'philosophical' approach to the questions of metaphysics, conceived as an approach different in its methods or objects from those already embraced by the empirical sciences. (1998, 4).

### Reply

Scientists can and should participate in scientifically responsible metaphysics. But it requires the input of philosophers, too. Philosophers can contribute something novel because we tend to approach questions differently than scientists do — from the top down, rather than the bottom up.<sup>63</sup> That is, we often begin by examining our overarching concepts and frameworks, rather than with particular pieces of data. Both approaches can produce valuable results. Moreover, philosophical training equips us with skills critical to the success of scientifically responsible metaphysics. Philosophers are trained in clear and rigorous argumentation and argument evaluation. *Contra* Ladyman and Ross (2007), metaphysicians are very good at error detection. As we saw in Chapter 1, while metaphysical claims are not subject to the same kinds of error as are

---

<sup>63</sup> Credit is due to Michael Devitt for this point.

empirical claims, they are subject to errors of reason, such as fallacies, failures of validity and cogency, vulnerability to counterexamples, and so forth. Much metaphysical literature consists of philosophers pointing out those errors. Scientific training focuses less on developing these types of error detection than does philosophical training. The result, when scientists attempt to do philosophy, can be disastrous. For instance, physicist Lawrence Krauss published a book (2012) claiming to show from considerations of physics how something could come from nothing. Famously, David Albert responded with a scathing review in the *New York Times* (2012), showing that the book was fundamentally conceptually confused, specifically regarding the concept of *nothing*. This goes to show that scientists who attempt to do philosophy sometimes lack the skills and training to do it well. See also Fodor (1998) and Rey (1985) on confusions about concepts among cognitive scientists.<sup>64</sup> So we shouldn't turn the project of scientifically responsible metaphysics entirely over to scientists. Rather, philosophers with suitable knowledge of science are vital participants in scientifically responsible metaphysics.

### 3. Objection from underdetermination

If the problem with free range metaphysics is that it is not sufficiently constrained, why should we think that science fares any better? Science, too, as Duhem (1914), Quine (1951) and van Fraassen (1980) show, is arguably underdetermined in various ways (see Laudan 1990b for disambiguations). In this section, I will understand *underdetermination* to mean that two or more mutually inconsistent theories are logically compatible with the available evidence and satisfy the constraints on theoretical content. In the case of scientific metaphysics, we have underdetermination *compounded*: observational evidence underdetermines the scientific theory

---

<sup>64</sup> I thank Michael Devitt for this example.

and the scientific theory underdetermines the metaphysics (see French 1998). So while the scientific constraint may secure epistemic warrant, it is not clear that it meets the robustness criterion.

### **Reply**

Indeed, current science is by no means a determinative constraint on the content of metaphysical theories. And there are indeed multiple levels of underdetermination underlying scientific metaphysics. But this is simply epistemic life. Underdetermination is pervasive; it arises as soon as one departs from one's evidence. Now, it is perfectly consistent for me to grant underdetermination. That is because to say that the constraints on free range metaphysics fail to be robust is *not* simply to say that they underdetermine the metaphysical theory. If it were, then a scientific constraint would not resolve the problem. It is to say, rather, that free range metaphysics is *unacceptably* underdetermined. There is a spectrum of underdetermination — some sets of evidence and theoretical constraints are compatible with greater numbers of rival theories than other sets. This is just to say, as I said in Chapter 2, that some theories are more constrained and some less. The stronger the constraints on theoretical content, the less leeway we have with that content, and the fewer rival theories satisfy those constraints. Thus, my claim that some levels of constraint are acceptable and some not translates into the claim that some levels of underdetermination are acceptable and some not. So I may be understood to have claimed that, from an epistemic point of view, free range metaphysics is unacceptably underdetermined, while scientific metaphysics is acceptably underdetermined.

#### **4. Objection from scientific anti-realism**

Suppose that we are independently convinced of scientific anti-realism. On that view, current science is not even approximately true. Our current scientific theories will be replaced over time with substantially different theories. If so, why should we base metaphysical theories on substantially false scientific ones? Wouldn't the substantial falsity of our current scientific theories contaminate our metaphysics? And wouldn't we thereby have a failure of epistemic warrant?

#### **Reply**

I reply that we need not build into scientifically responsible metaphysics the aim of tracking truth. Of course, it would be nice if it did. But it may be the case that that is too lofty an epistemic aim given our early and incomplete state of intellectual advancement. If scientific anti-realism is true, we may conceive of scientifically responsible metaphysics as an exercise in developing the most *justified* theories, rather than an exercise in immediately tracking truth. As Ladyman and Ross note, if our metaphysics is “at least motivated by our most careful science at time  $t$ , then it is the best metaphysics we can have at  $t$ ” (2007, 2). And seeking greatest justification is a perfectly respectable exercise — epistemically, we must do our best in the circumstances in which we find ourselves. However, for the purposes of this dissertation, I remain neutral on the issue of scientific realism.

#### **5. Objection from new cynicism**

A skeptical response to calls to naturalize metaphysics more broadly might take inspiration from what Susan Haack (2003) calls new cynicism. New cynics question the epistemic credentials of science. As Haack explains, some



radical sociologists, feminists, and multiculturalists, radical literary theorists, rhetoricians, and semiologists, and philosophers outside strictly philosophy-of-science circles... [insist] that the supposed ideal of honest inquiry, respect for evidence, concern for truth, is a kind of illusion, a smokescreen disguising the operations of power, politics, and rhetoric... Science is largely or wholly a matter of interests, social negotiation, or of myth-making, the production of inscriptions or narratives; not only does it have no peculiar epistemic authority and no uniquely rational method, but it is really... just politics. (2003, 20-21)

On this view, science is chock-full of ideology and political agendas and does not contain ‘absolute truth’. We might relatedly mention the role of big money — of corporate and governmental bodies whose funds drive the course of research. If so, then we shouldn’t assign science epistemic distinction. We shouldn’t trust it to inform us about the world and we shouldn’t trust a metaphysics that hinges on it.

## **Reply**

The value of this ‘new cynical’ view is its reminder that no form of inquiry proceeds in a vacuum. To the extent that all forms of inquiry operate within a socio-political context, all of them are shaped by socio-political forces — *inter alia*, by ideology, politics, power, and big money. It is important to remember that and interesting to consider the specific ways that those forces manifest in different forms of inquiry. The danger of the view is in the conclusion that science is *just* politics and in the accompanying denial of its epistemic distinction. It doesn’t follow from something having political dimensions or being shaped by political forces that it reduces somehow to politics. And of course, the epistemic distinction we assign to science isn’t premised on its being a-political. It is premised, rather, on its great empirical success with regard to explanation, novel

prediction, and the facilitation of technological innovation, successful intervention, and so on. The socio-political dimensions of science are not at odds with its epistemic distinction and the great success of science should limit our cynicism. For fuller replies to those who have challenged the epistemic credentials of science, see Boghossian (2006), Devitt (1997), Gross and Levitt (1998), Haack (2003), Fox (1988), Kitcher (1982), Koertge (1997), Laudan (1981), and Pennock (2000).

## **6. Objection from relevance**

One might think that science just does not bear on many or most traditional metaphysical questions. For instance, one might ask, what does science have to say about whether numbers, forms, universals or other abstract objects exist? These are matters that science simply does not speak to. And that is not contingent — *in principle*, science cannot inform us about those objects, since they are both unobservable and causally inert.

## **Reply**

Science cannot inform us directly about certain traditional objects of metaphysical theorizing. However, there may be scientifically constrained arguments, such as indispensability arguments, that motivate us to posit certain abstracta. For instance, one might commit to numbers because physics quantifies over numbers. Furthermore, I haven't claimed that *every thesis* or *every posit* needs direct support from science. If claims about abstract entities perform some needed explanatory work in a scientifically-informed worldview *and* respect the holism-related rules I put forward at the end of §II.2 above, very well. So science may not need to bear directly on a topic for it to figure into a scientifically respectable metaphysics. If, on the other hand, our claims about

such topics don't perform needed explanatory work and respect my rules for implementing holism, then they remain the proper subject-matter of free range metaphysics.

### **Sum**

I have argued for two theses: first, that free range metaphysics shouldn't be thought to produce epistemically justified theories of the world, and second, that metaphysicians who care to produce such theories should engage in scientifically responsible metaphysics. I argued that free range metaphysics is epistemically inadequate because it is not robustly constrained and its methods fail to secure epistemic warrant. I then argued that scientifically responsible metaphysics is not similarly inadequate because scientific theses are systematically confirmed and disconfirmed. By virtue of disconfirmation, the scientific constraint is robust, and by virtue of confirmation, it secures epistemic warrant.

## CHAPTER 4

### Case Studies

This chapter will perform case studies of two literatures in metaphysics, examining epistemological commitments and practices common to each, as well as the prospect of making them scientifically responsible. I have purposely selected two difficult cases for the would-be scientifically responsible metaphysician — cases where science seems, perhaps of necessity, irrelevant to the metaphysical claims one may want to make. The metaphysical topics I will examine are not obviously apt to be made scientifically responsible: 1) modal metaphysics and 2) grounding. While metaphysical topics that are *genuinely* not apt to be made scientifically responsible can be relegated to free range metaphysics, we must work to determine which topics are *in fact* apt or not, bearing in mind that as science and free range metaphysics advance, new and unforeseen points of contact may emerge between them. So the scientifically responsible metaphysician should continually look for footholds to be gained, where free range metaphysics can be reined in and made scientifically responsible. Where we find no present footholds, we wait in hope for future ones.

In each case study below, I begin by describing the state of the literature. I examine the explicit epistemology that is meant to capture and govern theorizing in the given metaphysical domain, as well as the first-order practices of metaphysicians working within it. My aim is to assess the potential for scientific responsibility and to identify potential points of contact between the given subject-matter and science. I then make a number of concrete recommendations for how to make work in those domains scientifically responsible. The chapter helps flesh out what scientifically responsibility looks like in these concrete domains.

## Case I: Modal Metaphysics

Case I deals with modal metaphysics. Modal metaphysics is problematic for a would-be scientifically responsible metaphysician because *prima facie*, actuality does not acquaint us with (non-actual) possibility or necessity. In this first case study, I will survey rationalist, Williamsonian, and empiricist metaphysics and epistemology of modality. Here I will not separate out epistemology from metaphysics, since the two are closely intertwined. I will find that none of the extant literature reveals avenues to making the metaphysics of modality scientifically responsible. However, I will argue that modal metaphysics can be made scientifically responsible by taking current science as an evidence-base for the justification of modal claims and as a model of good modal reasoning. I will argue that in addition to being better constrained and securing greater epistemic warrant than its free range alternatives, scientifically responsible modal metaphysics allows for greater resolution of disagreement and accounts for modal error.

### Preliminaries

I construe *modal metaphysics* broadly, such that it investigates not just the nature of modality, but also what the modal facts are. I take modal metaphysics to be the study of *metaphysical* possibility and necessity. For my purposes, *metaphysical possibility* is possibility not holding fixed the laws of physics.<sup>65</sup> It is, as van Inwagen put it, “Possible *tout court*. Possible *simpliciter*. Possible *period*... I mean possibility without qualification” (1998 72). We have a metaphysical possibility when some state of affairs could have been otherwise. *Metaphysical necessity* is not merely the local necessity with which one set of conditions physically necessitates

---

<sup>65</sup> Note that this is not uncontroversial as a gloss on metaphysical modality, since some philosophers have argued that metaphysical modality is co-extensive with physical modality (Ellis 2002, Mackie 2006, Shoemaker 1998).

another, but something broader — something that holds across all possible worlds. We have a metaphysical necessity when some state of affairs could not have been otherwise.

## **I. The State of Modal Metaphysics, Epistemology and Practice**

### **I.1 Modal Rationalism**

Traditionally, modal metaphysics assumes a rationalist epistemology. *Modal rationalism* is the view that knowledge of modal truth is gained primarily *a priori*. As I have described it, the view is descriptive. However, it need not be a purely descriptive view. *Prima facie*, modal rationalism is the *proper* epistemology for modal metaphysics. That is because the domain of modal metaphysics is precisely *not* the empirically discoverable domain of the actual. It is not immediately clear how the empirical — and science in particular — could acquaint us with this non-actual domain. As Hale puts the point:

experience — roughly, sense-perception and introspection, together with what we can infer from their deliverances — leaves us in the dark about what *might* be... the interesting question concerns knowledge of *unrealized* possibilities... It is precisely because possibilities may go unrealized that experience cannot teach us what must be so. (2002, 1)

Seemingly what *might* or *must* be is not revealed to us by experience, whether scientifically or otherwise. For that reason, the modal rationalist thinks there must be some non-empirical mode of epistemic access to modal truths. Conceivability and understanding are two candidate modes of epistemic access.<sup>66</sup>

---

<sup>66</sup> Devitt (2010) argues that such accounts put the cart before the horse, in the sense that they attempt to give an account of how we get knowledge of metaphysical facts without first determining what the facts are.

### I.1.1 Conceivability

According to conceivability accounts of modal knowledge, thought is our guide to modal truth. Conceivability accounts aim to capture the Humean maxim that “whatever the mind clearly conceives, includes the idea of possible existence, or in other words, that nothing we imagine is absolutely impossible” (Hume qtd. in Yablo 1993, 1). Some think the relation between conceivability and modal truth is an entailment relation (Chalmers 2002). For instance, on Chalmers’ view, if some statement *S* is conceivable in a certain way (we may omit the details for our purposes here), then there is a metaphysically possible world that satisfies *S* when considered as actual (2002). Note that this is a metaphysical claim. Others think that conceivability *evidences* rather than entails modal truth (Menzies 1998, Yablo 1993). Note that this is an epistemic claim. Further, some conceivability accounts are *non-epistemic* in the sense that they are not concerned with actual knowers in particular epistemic circumstances, but with ideal conceivers (Chalmers 2002, Geirsson 2005). In contrast, *epistemic* conceivability accounts posit conditions for knowers in actual epistemic circumstances having modal knowledge (Worley 2003, Yablo 1993). On such views, “whether or not something is conceivable for a thinker depends on what the thinker knows or believes, or what concepts or modes of presentation he has available or is using to think about the situation” (Worley qtd. in Roca-Royes 2011, 24).

For a specific example of a conceivability view of modal knowledge, take Yablo’s (1993) account. According to Yablo, “what is conceivable is typically possible, and... *p*’s conceivability justifies one in believing that possibly *p*” (1993, 13). Yablo explains, “Just as someone who perceives that *p* enjoys the appearance that *p* is *true*, whoever finds *p* conceivable enjoys something worth describing as the appearance that it is *possible*” (1993, 5). In his view, *p* is conceivable for me if “I can imagine a world that I take to verify *p*” (1993, 29). Although conceivability is not, in

his view, an infallible guide to possibility, it is still justificatory: “*probably*, if my evidence holds, then so does my conclusion” (Yablo 1993, 17). Yablo’s claim is not that conceivability is the *best* guide to possibility, but rather that it is the *only* guide: “if there is a seriously alternative basis for possibility theses, philosophers have not discovered it” (1993, 2). The maxim that conceivability is a guide to possibility is, he thinks, “entrenched, perhaps even indispensable” (1993, 2).

Chalmers (1996) provides an example of the conceivability-based epistemology of modality operating at the first order — that is, in the *practice* of modal metaphysics. He asks us to consider the logical possibility of zombies, where a zombie is “just something physically identical to me, but which has no conscious experience—all is dark inside” (1996, 96). He says that the logical possibility of these zombies “seems... obvious” (1996, 96). Having described zombies, he argues as follows:

While this is probably empirically impossible, it certainly seems that a coherent situation is described; I can discern no contradiction in the description. In some ways an assertion of this logical possibility comes down to a brute intuition... Almost everybody, it seems to me, is capable of conceiving of this possibility. (1996, 96)

In his view, the conceivability of zombies secures the claim of their logical possibility. Elsewhere in the book, Chalmers argues that “the metaphysically possible worlds *are* just the logically possible worlds” (my emphasis, 1996, 38). So Chalmers gives us a conceivability argument for the metaphysical possibility of zombies.

### **I.1.2 Understanding**

In contrast, some rationalist epistemologies of modality take understanding to be the primary mode of access to modal truth. For instance, Peacocke (1999) articulates an



understanding-based account, according to which modal knowledge hinges on the understanding of modal notions, which carries with it implicit knowledge of the principles of possibility. In his view, modal knowledge consists in the proper use of those implicitly known principles (1999, 162). Moreover, Bealer (2002) gives an account according to which modal knowledge is a function of the determinate understanding of concepts. A subject determinately understands a concept if and only if she has “natural propositional attitudes toward propositions that have that concept as a constituent content” (2002, 102) *and* “does not do this with misunderstanding or incomplete understanding or merely by virtue of satisfying our attribution practices or in any other such manner” (2002, 102). Determinate understanding of concepts, together with good “cognitive conditions (intelligence, etc.)”, yields truth-tracking modal intuitions (2002, 103). If a subject’s intuitions were not truth-tracking, “the right thing to say would be that either [she] does not really understand one or more of the concepts involved, or her cognitive conditions are not really good” (2002, 103). Bealer gives the following example: a person with good cognitive conditions determinately understands the concept *multigon*. Suppose that to be a multigon is to be “a closed, straight-sided plane figure” (2002, 103). Her determinate understanding of the concept generates in her the intuition that “it is possible for a triangle or a rectangle to be a multigon” and her intuition is truth tracking (2002, 103). Peacocke and Bealer both articulate views according to which modal knowledge is a function of conceptual understanding.<sup>67</sup>

Let me give an example of conceptual understanding putatively yielding modal knowledge in first-order metaphysics, broadly construed: the ontological argument. Some versions of the ontological argument are premised on there being a *concept* of God that ascribes to him attributes

---

<sup>67</sup> Note that in this respect, Peacocke and Bealer fall in a tradition that tried to make analyticity a route to *a priori* knowledge of necessity (see Devitt 2011).

like *greatest conceivable greatness* (Anselm 1078), *maximal excellence* (Plantinga 1974), or *supreme perfection* (Descartes 1637, Leibniz 1709). Some versions ultimately conclude that God exists *necessarily* — that is, exists in all possible worlds (e.g. Anselm’s second ontological argument, on interpretations such as Malcolm’s 1960). Here we have a metaphysical argument that derives a modal conclusion from purported conceptual understanding.

## **Taking Stock**

Conceivability and understanding-based accounts of modal knowledge frequently assign intuition a prominent evidentiary role. For instance, Bealer thinks intuition is a basic form of truth-tracking evidence given the right kind of conceptual understanding (2002, 102). Moreover, rationalist modal metaphysics frequently proceeds on the basis of intuition. The metaphysician uses thought experiments to draw attention to intuitions about phenomena, for the purposes of establishing particular modal conclusions. Examples include Putnam’s (1973) twin earth thought experiment and Chalmers’ (1996) zombie thought experiment. In Chapter 3, I argued for the epistemic dubiousness of such intuitions. Constrained primarily by intuition and absent contact with science, any modal metaphysics that adopts or assumes a rationalist epistemology counts as free range metaphysics. Therefore, the criticisms I levelled against free range metaphysics in Chapter 3 apply to rationalist modal metaphysics. Thus far, the examination of the literature hasn’t revealed any promising avenues for making modal metaphysics scientifically responsible.

### **I.1.3 Further Difficulties**

Though my arguments in Chapter 3 are sufficient to undermine the epistemic credentials of free range modal metaphysics, Yablo (1993) points out further difficulties facing modal

rationalists. First, for those modal rationalists who assign an evidentiary role to modal intuition, disagreement often bottoms out in the dull thud of clashing intuitions (Chalmers' 1996 zombie thought experiment is a particularly divisive example). As Yablo puts it, "more often than not we *still* bog down" in "clashes of modal intuition" (1993, 40). When we get these brute disagreements, "the faith that there is anything genuinely at issue can indeed become strained" (Yablo 1993, 38). Unless we find some grounds for reproaching the opposing view, it seems the only option is "to insist on there being 'facts of the matter' that only oneself and one's coreligionists are privy to" (Yablo 1993, 38). And at that point, we get "something of a credibility problem" (Yablo 1993, 38). If we are to maintain the view that some modal claims are false and that there is a consistent set of modal facts with which a responsibly formed theory can acquaint us, then modal disagreement should not be brute. Rather, there should be some way of gaining traction on our disagreements.

Moreover, we need some account of why and when modal error occurs, and modal rationalism does not provide such an account. Yablo tells us "[m]odal error is a fact of life, and although perceptual error is too, our firmer grip on its etiology allows us to feel less the helpless victim than in the modal case" (1993, 32). We can usually assume that perceptual appearances are more-or-less accurate, because "[m]isperception is something that we know how to guard against, detect when it occurs, and explain away as arising out of determinate cognitive lapses" (Yablo 1993, 32). But this is not the case with modal seemings. Yablo explains, "there is nothing remotely comparable for conceivability" and this is an indication of "our relative backwardness on the subject of modal error" (1993, 32-3). The point stands not just for conceivability, but for modal intuition generally. Yablo concludes that "[u]ntil our imaginative excesses are brought under something *like* the epistemological control we have in other areas, we modalize with right, perhaps, but without conviction" (1993, 33). I say that we modalize with *neither* right nor conviction. Truly

satisfactory epistemologies of modality should be able to account for modal error, and modal rationalism does not.<sup>68</sup>

## **I.2 Williamson's Counterfactual Account**

Williamson (2007) articulates an alternative epistemology of modality, according to which the capacity to modalize is a product of our ability to think counterfactually. His view lacks the normative dimensions that I attribute to modal rationalism and will attribute to modal empiricism, and is instead purely descriptive. He considers his view neither rationalist nor empiricist, since it characterizes modal knowledge neither as purely *a priori*, nor as purely *a posteriori* — a distinction he thinks “obscures... significant epistemic patterns” (2007, 169). In his view, while modal knowledge is gained from the armchair, it nevertheless “[does] not fit the stereotype of the *a priori*, because the contribution of experience [is] far more than enabling” (2007, 169). The role of experience is more than enabling because of the centrality of counterfactual thinking to modal knowledge, together with the fact that counterfactual thought is “deeply integrated into our empirical thought in general” (2007, 141). In Williamson's view, “the ordinary cognitive capacity to handle counterfactual conditionals carries with it the cognitive capacity to handle metaphysical modality” (2007, 136). So the story of modal knowledge is really the story of mundane counterfactual knowledge. That story is developmental and evolutionarily un-mysterious. As I grow and learn, experience “condition[s] me into patterns of expectation which are called on in my assessment of ordinary counterfactual conditionals” (2007, 168). We gradually accrue a background of empirical knowledge that informs everyday counterfactual thoughts, like “If the

---

<sup>68</sup> This objection to modal rationalism may count against *a priori* knowledge more broadly. I thank Michael Devitt for the point.

bush had not been there, the rock would have ended in the lake” (2007, 142). Our counterfactual thinking is often the operation of our imagination, “radically informed and disciplined” by the empirical background of beliefs and an accompanying folk physics (2007, 143). Imagination figures in frequently, but not always, since “imaginative simulation is neither always necessary nor always sufficient” for the evaluation of counterfactuals (2007, 152). In sum, according to Williamson’s modal epistemology, much modal knowledge is to be accounted for in terms of counterfactual knowledge.

Examples of putative modal knowledge fitting Williamson’s account must be suitably mundane. That is because, since “we may well be more reliable in evaluating counterfactuals whose antecedents involve small departures from the actual world”, we may likewise be “more reliable in evaluating the possibility of everyday scenarios than of ‘far-out’ ones” (2007, 164). So our example should be something like this: I know that it’s possible for me to leave my desk right now because, by virtue of an empirically-informed background of belief, I can counterfactually imagine doing so. If Williamson’s account is descriptively correct, it may suggest that modal metaphysicians who speculate about more ‘far-out’ possibilities stretch their capacity to modalize beyond the limits of everyday counterfactual thinking — and that may undermine their claims to knowledge.

### **Taking Stock**

There may be a lack of fit between Williamson’s account of modal reasoning and the activities of modal metaphysicians. Williamson’s account may well describe everyday modal reasoning, but it may not capture the more esoteric explorations of distant possible worlds that metaphysicians engage in. So his account may not be particularly germane for my purposes here.

But suppose that Williamson’s account adequately describes some reasoning on the part of modal metaphysicians. Notwithstanding the empirical credentials of their reasoning, it may still be free range. That is because Williamson-style modal (that is, counterfactual) reasoning need not operate against a scientifically-informed background of belief. It operates against a folk theory. So Williamson’s account does not help us to see what scientifically responsible modal metaphysics looks like.<sup>69</sup>

### I.3 Modal Empiricism

Modal empiricism is developing at the forefront of current modal epistemology as an exciting new alternative to the traditional rationalist view (see Fischer and Leon 2017). For my purposes, *modal empiricism* is the view that modal truths are known to some extent *a posteriori*. Such views come in varying strengths: modest versions say that *some* modal truths are known *a posteriori* (Roca-Royes 2007), while the stronger versions say that *many* of them are (Bueno and Shalkowski 2014).<sup>70</sup> The view can also be characterized normatively, so that modal empiricism is the *proper* route to knowledge of some or many modal truths.

*Prima facie* motivations for modal empiricism include the following. First, “actuality is often the best argument for possibility” (Williamson 2007, 164). That is because the actual *demonstrates* the possible. However, it demonstrates (presumably) only a small subset of the possible. So this motivation does not get us far: it only pushes us to accept empiricism about a small subset of modal

---

<sup>69</sup> However, plausibly, our background knowledge *could* be filled out and made precise by science learning, in a way that would better enable counterfactual imagination. I will make a similar suggestion in section II.1 below.

<sup>70</sup> Note that, so understood, modal empiricism is not antithetical to versions of modal rationalism that grant some *a posteriori* modal knowledge. Roca-Royes explains: “for some of the knowable modal facts, an empiricist story is to be told, whereas a rationalist one may be adequate for the rest” (2007, 119).

truths. A further *prima facie* motivation is that science acquaints us with important modal facts. Take Kripke's (1980) famous examples of *a posteriori*-known essential properties: water is essentially H<sub>2</sub>O, gold is essentially atomic number 79. We discovered these properties empirically. However, post Kripke, most modal rationalists acknowledge these *a posteriori*-known modal facts and simply relegate them to the minority (Vaidya 2015, §2). Moreover, a modal rationalist may respond that in the Kripke examples, though science acquaints us with the relevant *properties*, it does not acquaint us with their *modal profiles*.<sup>71</sup> That is, it does not reveal that the properties are essential. We know that they are essential by virtue of our judgment, not empirically. So the *prima facie* motivations don't sufficiently motivate modal empiricism. A better motivation for modal empiricism would show that some or many modal truths are known empirically.

Several modal empiricists try to show just that. Rebecca Hanrahan argues for a version of modal empiricism that hinges on the following analogy: "the imagination is to the possible as perception is to the actual" (2009, 282). That is, imagination fallibly justifies beliefs about what is possible. In her view, imagination is to be understood as Hume understood it:

the images produced by our imagination are constructed in some way out of elements of what we have previously perceived. So... the images produced by the imagination could have been produced through the workings (standard or otherwise) of our sensory faculties.

(2009, 293)

So the view is empiricist to the extent that it takes imagination to be *a posteriori* – to be dependent on and derivative of experience. Hanrahan gives the following example: while cooking, "a series of images came to me... of Walter accidentally slicing through one of his fingers... and they were for the imagination quite forceful and vivacious" (2009, 292). That imaginative experience

---

<sup>71</sup> I am grateful to Gary Ostertag for pointing this out to me.

justifies belief that the proposition *Walter has cut off one of his fingers* is possibly true (2009, 292). According to this empiricist epistemology of modality, imagination, understood as an *a posteriori* form of evidence, is our guide to possibility. Alternatively, Carrie Jenkins (2010) argues that conceivability is a guide to modal truth, but that conceivability is a function of empirically-gained conceptual competence. She explains, “the senses may ground modal knowledge by providing what I call epistemic grounding for our concepts, which concepts (help to) determine what we can and cannot conceive of, which in turn guides our modal beliefs” (2010, 255).<sup>72</sup> Margot Strohminger argues that sense perception acquaints us with many non-actual possibilities, so that “‘I see that I can reach the mug’, ‘I see that I can climb the tree’, and ‘I see that I can catch the ball’” (2015, 367). The literature on modal empiricism is at an early, exploratory stage of development. Each view remains to be worked out and it remains to be seen which of them will emerge as serious theoretical contenders.

### **Taking Stock**

Modal empiricism takes us a step toward scientifically responsible modal metaphysics, since the scientific is a species of the empirical. Modal empiricism reconceives of, or demonstrates alternatives to, the more mysterious *a priori* methods of the modal rationalist. But none of the versions of modal empiricism we have seen point to a role for science. None of them show how scientific knowledge might figure into or ground modal knowledge. Like the other epistemologies of modality that I have considered, my examination of modal empiricism has failed to illuminate how modal metaphysics might be made scientifically responsible.

---

<sup>72</sup> See Devitt’s (2014) response.



## II. Scientifically Responsible Modal Metaphysics

In contrast, Daniel Nolan (2017) gestures toward a naturalized modal epistemology. He suggests that a naturalistic modal epistemology is one that uses scientific techniques, including the techniques of psychology, to investigate how we *do* come to know modal facts. However, for my purposes here, I am not interested in how science can illuminate *how we know* modal facts, but in how science can illuminate *the modal facts*. That is, I wish not to articulate a scientifically responsible modal epistemology, but to articulate an epistemology, or methodology, for scientifically responsible modal metaphysics. The former would use science to investigate modal knowledge; the latter would say what role science should play in the justification of modal metaphysical claims. In this section, I will develop that methodology. I will spell out what I mean by *scientifically responsible modal metaphysics*, and in particular, the relevance that science has to modality.

While Nolan's (2017) modal epistemology is not relevant for my methodological purposes, he does make a suggestion that I think is important for those purposes. He claims that empirically informed work on conditional matters, dispositional matters, and ability claims can yield generalizations explainable by a systematic theory. We may find answers to less empirically tractable modal questions by appeal to that empirically informed theory. The suggestion is similar to Callender's (2011) claim that "*we should take possibilities and necessities only as seriously as the theories that generate them*" and that the theories that generate them should be comprehensive and empirically virtuous (original emphasis, 2011, 44). Nolan and Callender both make the holist point that any claim we take seriously should figure into an overall theory supported by our best available evidence. So a methodological principle for scientifically responsible modal metaphysics

might be: take seriously the modal claims that figure into scientifically-informed worldviews, where ‘figure into’ could mean ‘are explainable by’ or ‘are generated by’.

But in my view, science can bear on and guide modal metaphysics more robustly. One of Nolan’s suggestions is promising here. He argues that we may gain insight into how to support modal claims by studying the manner in which scientific models produce modal information (2017). I will argue something similar below. In my view, a *scientifically responsible modal metaphysics* justifies its modal claims inductively on the basis of scientific evidence or models itself after the modal reasoning that goes on in science. But for scientific responsibility to be appropriate in the context of modal metaphysics, science must speak to modal truth, and it does not *obviously* do so. However, let me now explain its relevance.

## **II.1 Science and Modal Truth**

How does science bear on modality? Let me start with the relevance of science to metaphysical possibility. Other than acquainting us with actual demonstrations of the possible, science has two dimensions of relevance to metaphysical possibility. First, science facilitates knowledge of physical possibility by acquainting us with physical laws. Given that physical states of affairs are either consistent or inconsistent with the laws of physics, they are either physically possible or not. And, since physical possibility is a subset of metaphysical possibility, knowledge of physical possibility is knowledge of metaphysical possibility (note that I am so far claiming nothing about metaphysical necessity). For all those metaphysically possible worlds with physical laws similar to our own, science facilitates knowledge of them.

The second dimension along which science is relevant to metaphysical possibility is less direct. Science aids us by being a model of the responsible exploration of possibility space. The

exploration of possibility space is integral to scientific practice. When scientists predict actual and counterfactual outcomes and test competing hypotheses, they consider an array of physically possible outcomes. To facilitate description, explanation, and prediction of phenomena, scientists frequently posit laws, mechanisms, and models. See for instance the following uses of causal modelling: Bouzat (2014); Chennu, Noreika, Gueorguiev, Shtyrov, Bekinschtein, Henson (2016); Giret, Kornfeld, Surya, and Hahnloser (2014); Park and Kim (2015); Schneider, White, and Paul (1998). Scientists' knowledge of natural laws and their tools for modelling observed phenomena motivate inferences to certain physical possibilities. Scientists determine which physical possibilities are consistent with or likely given the laws and the models. So scientific reasoning routinely includes modal reasoning about physical possibility, and scientists have developed tools and techniques that aid in that reasoning. Reasoning about broader *metaphysical* possibilities is an activity similar in kind. That is, when we reason about non-physical possibility, we do the same *kind* of thing as when we reason about physical possibility. Namely, we check to see which propositions are consistent with a certain set of propositions that we hold fixed.<sup>73</sup> When we reason about non-physical possibility, the difference is that we vary the propositions we hold fixed — we vary laws against which we judge things to be possible or not. Knowing how to reason well about physical possibility can illuminate how to reason well about broader metaphysical possibilities. That is the role that science should play with respect to our explorations of non-physical possibility: it should be taken as a model of good modal reasoning.

As for the relevance of science to necessity, science reveals some necessities when it reveals laws of nature. I have in mind dynamical laws governing, say, motion and transference of energy. If we take those laws to be metaphysically necessary (as do the *necessitarians*, including

---

<sup>73</sup> I'll say more about this in §II.2 below.

Ellis 2001, Swoyer 1982, Shoemaker 1980), then when science posits laws, it acquaints us with metaphysical necessities. But even if we think physical laws are not metaphysically necessary, science illuminates metaphysical necessity in the following ways. It acquaints us with candidate essential properties (e.g. it acquainted us with the chemical composition of water, which we then judged to be essential) and numerical identities. Moreover, if, when we vary the physical laws, we find that certain modal states of affairs hold fixed across variations, then we will have acquired inductive evidence of necessary possibilities.<sup>74</sup> In other words, the more we find that  $\diamond p$  holds across worlds, the more inductive evidence we have of the metaphysical necessity ( $\Box_m$ ) of  $\diamond p$  — that is, of  $\Box_m \diamond p$ . Those are the metaphysical necessities available to a scientifically responsible modal metaphysics.

A helpful way of capturing my claims thus far is as follows. In relation to modal truth, science has:

- 1) *direct bearing* where we hold fixed the laws of physics
- 2) *indirect bearing* where we vary the laws of physics

Where the bearing is direct, science more immediately furnishes us with modal facts. Where it is indirect, science provides a model for the responsible exploration of modal space.

Let's consider some examples. What sorts of modal claims would a scientifically responsible modal metaphysics commit to? Let's start with the necessities. First, if we hold the controversial view that physical laws are metaphysically necessary, and if laws like the first and second law of thermodynamics are genuine laws of nature, then those laws are metaphysically necessary (but we need not hold the controversial view). Moreover, scientifically responsible

---

<sup>74</sup> Credit is due to Paul Simard-Smith for his suggestion that explorations of metaphysical possibility might yield inductive evidence of certain necessities.

modal metaphysics recognizes the (partial) role of science in discovering certain essential properties, and so commits to the necessity of *water is H<sub>2</sub>O* (though as I pointed out above, arguably science does not get us all the way to the modal claim here). Further, being committed to empirically discovered numerical identities, it commits us to *Hesperus is Phosphorus*. Now let's consider the possibilities. First, science commits us to physical possibilities. Competing scientific hypotheses often signal different physical possibilities. For instance, consider the following competing hypotheses: 1) selection favours greater virulence among parasites and pathogens; 2) selection favours intermediate virulence (see Lenski and May 1994). Both are physically possible. Now, you might ask: what about the non-physical possibilities? And what about controversial cases, like the swamp-man and zombie cases? To answer those questions, we require some fully worked-out models of modal reasoning in science. Developing and applying those models is outside the scope of this dissertation — but it is an important avenue of future research. My aim here is not to work out a scientifically responsible modal metaphysics, but to say what developing one would involve.

## II.2 'Further Difficulties' Resolved

For the reasons I gave in Chapter 3, scientifically responsible modal metaphysics lacks the epistemic deficiencies of free range modal metaphysics. But it also betters modal rationalism with respect to the 'further difficulties' I spelled out in section I.1.4 above. Let me explain how.

First, the scientific constraint allows for greater resolution of disagreement. Although disagreement persists in science, as it does elsewhere, science has a good deal more traction on disagreement than other forms of inquiry because of the empirical testability of many (or most) of its theses. In the modal metaphysics I am recommending, disagreement can have the following

objects: 1) the scientific facts and 2) the modal profile of the scientific facts. Disagreements about 1) can be settled by checking the available body of scientific data. Disagreements about 2) can be settled in the following ways. Suppose  $M$  is a proposition whose modal status we wish to evaluate. To check whether  $M$  is metaphysically necessary, we check the actual science. Is  $M$  a law? Is there sufficient evidence to make it a law? Is it a candidate essential property or numerical identity? We should also consider whether there is inductive evidence that  $M$  holds across all possible worlds. To check whether  $M$  is possible, we check first whether it is consistent with the physical laws. If not, we check whether it is consistent with some other set of laws. Let us say that we are evaluating  $M$ 's modal status at some distant possible world — so we are not holding fixed actual empirical observations or physical laws (*pace* the necessitarians). To establish  $M$ 's possibility, we determine whether  $M$  is consistent with all premises about that world. Note that we cannot determine  $M$ 's possibility holding *nothing* fixed. In order to be consistent,  $M$  must be *consistent with* something. So we must specify what is to be held fixed, e.g. the laws of classical logic, the physical laws pertaining to that world, a background of counterfactuals, etc.  $M$  is either consistent with these specifications or not. If it is, then we have determined  $M$ 's metaphysical possibility. On this view, modal disagreement does not bottom out in clashing intuitions about conceivability.

Moreover, scientifically responsible modal metaphysics has the resources to explain and guard against much modal error. That is because in the case of direct bearing, modal error will be the result of scientific error. Scientific errors will primarily be perceptual errors, measurement errors, and mathematical errors — the sorts of errors that we have a good handle on. The errors may also result from incomplete scientific evidence, for instance, in a case where we previously had not observed the conditions under which some regularity breaks down. That error can be explained with reference to the fallibility of inductive inference and can be guarded against with a

sufficiently strict standard for the establishment of a law. In general, when we modalize holding fixed a background including the physical laws, the more precise the specification of physical laws, the better the resulting modal claim. On this view, modal error could also be the result of poor imitation of the scientific model, which we could certainly improve. So compared with rationalist modal metaphysics, scientifically responsible modal metaphysics has greater resources to account for, identify, and resolve modal error.

### **III. Conclusion, Case I**

In this first case study, I surveyed modal epistemology and metaphysics. Modal rationalism, Williamson's counterfactual view, and modal empiricism each failed to illuminate how modal metaphysics might fruitfully engage with science. I explained how modal metaphysics could be made scientifically responsible: by taking science as an evidence-base for modal claims and a model of good modal reasoning. Scientifically responsible modal metaphysics does not permit modal speculation to go where it will; it reins in modal speculation by tethering it to our most epistemically distinguished and empirically successful form of inquiry. In addition, it betters its rationalist alternatives by permitting fewer irresolvable disagreements and accounting for modal error. So scientifically responsible modal metaphysics is possible, and we now have a sense of how to go about it.

### **Case II: Grounding**

Case II deals with the grounding literature. The grounding literature is *prima facie* not apt to be made scientifically responsible, since it rejects the Quinean view of metaphysics in favour of an Aristotelian one (Schaffer 2009) and explores ontological dependence relations seemingly alien

to science (Hofweber 2009). In this second case study, I will examine some of the explicit epistemological commitments and first-order practices found in the grounding literature. Inspired by Audi (2012) and Schaffer (2003, 2010, 2015), I will then synthesize a list of recommendations for how science may be put to use in theories of ground, including among other things: to help identify putative grounding relata, to show correlations among them, to demonstrate their non-identity, to provide a stock of explanatory patterns, and to identify candidate essential properties. I will conclude that theories of grounding can be scientifically responsible and we have a general framework for making them so. Lastly, I will consider an objection from Hofweber that the grounding literature gives itself autonomy from science. I will respond that Hofweber simply shows the *need* for making theories of ground scientifically responsible in the way I suggest.

## **Preliminaries**

To say that  $x$  *grounds*  $y$  is to say, roughly, that  $y$  metaphysically depends on  $x$ . For something to metaphysically depend on something else is for it to obtain or exist *in virtue of* the other (Fine 2012). The thing doing the grounding is thought to be more fundamental, ontologically or explanatorily, than the thing being grounded. Beyond that, it is difficult to say what grounding-talk is meant to track. In part, that's because the term 'grounding' is often said to be primitive (e.g. Schaffer 2009). We are meant to have an intuitive grasp of it. However, there is little agreement about the relata of the grounding relation, its formal properties, or its paradigm examples. The following are some putative examples of grounding. This first is from Fine (2012) and the rest are from Schaffer (2009):

- the fact that A and the fact that B ground the fact that A&B
- the entity ground its singleton



- natural features ground moral features
- sparse properties ground abundant properties
- truthmakers ground truths

The examples don't impress everyone; some grounding skeptics have argued that the grounding concept plays no unique conceptual role (Wilson 2014). Let's set those skeptical worries aside and assume, for the moment, that the grounding literature is onto something.

## **I. The State of the Metaphysics of Grounding, Epistemology and Practice**

### **I.1 Epistemology**

I will begin by looking at some *second-order* claims of grounding theorists — that is, some of their explicitly epistemological and methodological claims. Here I will search for recommendations that might indicate possible avenues for making theories of ground scientifically responsible. I will then turn to the first-order practices of grounding theorists — that is, their metaphysical theorizing — and assess whether the methods that emerge in practice reveal any such avenues.

#### **I.1.1 Fine**

Let us begin with Fine, who I take to be the father (Aristotle, the grandfather) of this grounding literature. Fine makes his epistemology explicit. As we saw in Chapter 3, he assigns a role to two sources of evidence: intuition and explanatory power. He says: “There are, I believe, two main sources of evidence for making judgments of ground. The first is intuitive. We appear to be in possession of a wealth of intuitions concerning what does or does not ground what” (2001, 21). For instance, intuition leads him to say that “what grounds the truth of a disjunction is the

truth of those of its disjuncts that are true” and “what grounds the occurrence of a compound event at a given time is the occurrence of its component events” (2001, 21-22). But he thinks that:

[w]e also have intuitions about a wide range of negative judgements (quite apart from modal considerations). It is implausible, for example, that what grounds facts about volume are facts about density and mass or that what grounds the truth that a given object is red is the fact that it is red or round and the fact that it is not round, even though the one logically follows from the others. (2001, 22)

So intuition guides us toward certain grounding claims and away from others. It is the basis for our stock of paradigm examples. The second source of evidence, he tells us, is “explanatory in character” (2001, 22). On Fine’s view, when you provide the ground of a proposition, you explain it. Given the tight connection between grounding and explanation, Fine thinks “a system of grounds may be appraised, in much the same way as any other explanatory scheme, on the basis of such considerations as simplicity, breadth, coherence, or non-circularity” (2001, 22). The most important of these considerations, in Fine’s view, is explanatory strength, which he characterizes as “the capacity to explain that which stands in need of explanation and would otherwise be left unexplained” (2001, 22). And one particular thing that stands in need of explanation is the presence of a “certain necessary connection between the propositions that are so grounded” (2001, 22). So according to Fine, the epistemology of grounding assigns an evidential role to intuition and to explanatory virtue.

That is not to say, however, that *both* forms of evidence must support *every* grounding claim one makes. We may, for instance, lack intuitions about many particular grounding claims, say, when the subject-matter becomes less familiar. In those cases, Fine says that individual grounding claims can receive support on the basis of “holistic considerations” (2001, 22). So the

recommended methodology is a sort of balancing act between case-by-case judgments and considerations of fit with respect to “a general pattern of explanation” (2001, 22).

As I argued in Chapter 3, the constraints that Fine places on theoretical content are entirely free range, and scientific considerations do not figure in. Grounding theories constructed in accordance with his method are entirely independent of science. So Fine’s epistemology does not illuminate how theories of ground may be made scientifically responsible.

### **I.1.2 Schaffer**

Schaffer (2009) also makes explicit a methodology for formulating theories of grounding. According to his Aristotelian conception of metaphysics, “The method of metaphysics is to deploy diagnostics for what is fundamental, together with diagnostics for grounding” (2009, 351). There are three diagnostics for what is fundamental: 1) minimal completeness, i.e. the posited fundamentalia provide a supervenience base for the actual world and include nothing superfluous; 2) metaphysical generality, i.e. the most general features of the collection of posited fundamentalia are necessary features; and 3) what Schaffer calls *empirical specifiability*, i.e. “the substances have a content informed by fundamental physics” (2009, 377).

The last is the most relevant for my purposes here. He explains: “The content of an inventory [of fundamental substances] is its most specific features, and the content is empirically specifiable iff these features fit those found in fundamental physics” (2009, 377). It is important to note that while Schaffer characterizes this diagnostic as *empirical*, it is not merely empirical, but *scientific*. The distinction matters in this context because scientifically responsible metaphysics must be engaged with *science* in particular; merely being empirical is not sufficient. At any rate, I take Schaffer to mean that the fundamentalia posited by the metaphysics (call them *MP* for

metaphysical posits) must ‘fit’ the ones posited by fundamental physics (call them *SP* for scientific posits). This notion of ‘fit’ needs to be cashed out. Its role might be to require 1) that MP be logically consistent with SP, 2) that MP include nothing over and above the content of SP, or 3) that MP include all and only the items in SP. But none of these requirements secures scientific responsibility, in that none of them requires consilience with science to be the result of a conscientious and painstaking engagement with science. That is, 1-3 may occur by happy accident. To borrow Wilson’s (2006) phrase again, the empirical diagnostic appears to mention science without mandating its use. One might object that the diagnostics here are likely meant only as after-the-fact checks on theoretical results, rather than as guides to the formulation and justification of the metaphysical theory. I grant the point. But then the latter aspect of the methodology is missing. The point stands that the explicit methodology is not scientifically responsible, in that it contains no requirement for the metaphysical theory to be formed *on the basis of* science.

According to Schaffer, the diagnostics for the grounding relations include: 1) permissiveness, i.e. the grounding relations generate abundantly many entities (2009, 377) and 2) abstraction, i.e. the relations are not additions to the world, but merely ways of specifying aspects of it that are “implicitly present from the start” (2009, 378). Neither of these diagnostics even mentions science, so clearly neither secures scientific responsibility. In sum: the diagnostics for fundamentality and for grounding fail to secure scientific responsibility. Although Schaffer mentions science in his formulation of the empirical diagnostic, the requirement of fit does not ensure that any given metaphysical theory is formed in the right way, i.e. on the basis of science.

### I.1.3 Audi

Now I turn to Audi's epistemological claims, which are more promising as regards scientific responsibility. He agrees with Fine about the deep connection between grounding and explanation. He says the reason for countenancing a grounding relation at all is that doing so is "indispensable to certain important explanations" (Audi 2012, 104). And further, the justification of particular grounding claims hinges on our beliefs about explanation:

It might also be that claims about grounding are justified in a process of reflective equilibrium. We have beliefs about particular explanatory connections and we have general beliefs about the rules of explanation. Neither is indefeasible, but they can be mutually reinforcing. (2012, 117).

Audi suggests (following Daly 2012) that grounding claims can be backed empirically, albeit indirectly:

it is not clear that claims about grounding are insulated from empirical confirmation or disconfirmation. As Daly points out, facts about grounding might be indirectly empirically confirmable by having a place in a well-confirmed theory, a theory confirmed through its empirical consequences. (2012, 116)

So like Fine, Audi thinks that holistic considerations can figure into the justification of particular grounding claims; unlike Fine, however, he explicitly mentions the possibility of empirical support. As in the discussion of Schaffer above, I note that where Audi says *empirical*, we should read him as saying *scientific*. He articulates the step-by-step, scientifically-informed process by which we may be justified in believing particular grounding claims. He takes as an example the following grounding claim: *something's being an electron grounds its having negative charge*. He then says that the process of justification for that grounding claim is as follows:

1. We have contact with electrons in a way that enables us to refer to the property of *being an electron*...
2. We discover a disposition associated with being an electron, the power to repel certain things, and we call this power *negative charge*.
3. We discover that this disposition is associated with other things besides electrons, such as down quarks.
4. We may conclude from this that being an electron is not the very same thing as having negative charge.
5. We find reason to think that things' having negative charge is explained in terms of other facts about them (not the same fact in the case of electrons as in the case of down quarks).
6. These cannot be causal explanations.
7. We conclude that being an electron grounds having negative charge. (2012, 117)

While Audi does not generalize from the example, we may do so on his behalf. Here is the process, as I understand it, in a more schematic form:

1. We might have empirical reasons for positing entities of type  $x$ . So there is a property  $y$  of *being an  $x$* .
2. Entities of type  $x$  also have some empirically discoverable property  $\phi$ .
3. We discover that  $\phi$  is realizable by entities of types other than  $x$ ,
4. Therefore (from 1-3), we discover that  $\phi$  is not  $y$ .
5. We find reason to think that in general, the members of some entity type having  $\phi$  is explained in terms of facts particular to that entity type.
6. These explanations cannot be causal explanations.

7. So we conclude that in the case of some particular entity of type  $x$ , property  $y$  grounds its having property  $\varphi$ .

According to Audi, “this account is epistemically hybrid”, in that some of the steps are clearly empirical (read: informed by science) and some clearly *a priori* (2012, 117). In my view, the *a priori* step is the implicit assumption of Leibniz’s Law required to make the inference to step 4 valid. At any rate, Audi comments:

it would be oversimple to say only that grounding claims are a priori, or for that matter to say only that they are empirical. Like most substantive philosophical claims, their support derives in part from experience and in part from reflection on the information that experience provides. (Audi 2012, 117)

This is perfectly acceptable, and any methodology for substantive philosophy that is worth its salt will be hybrid in this way.

At any rate, Audi makes it clear that scientific evidence can help back grounding claims by allowing us to identify putative grounding relata (i.e. negative charge in the case of electrons), by showing correlations among those relata, by demonstrating non-identity of relata, and perhaps by providing a stock of explanatory patterns. Since Audi’s justificatory process involves bringing *scientific evidence* to bear in the justification of grounding claims, it is scientifically responsible in my sense of the term.

Nevertheless, Audi’s method will not figure as a whole into my recommendations. Let me explain why. First, the schema does not apply in cases where  $y$  and  $\varphi$  are essentially connected — e.g. *being water* and *having chemical composition H<sub>2</sub>O* — since step 3 will not be satisfied in those cases. The non-applicability of the schema in such cases is surprising, since Audi claims that being essentially connected is necessary for standing in a grounding relationship (2012, 108).

Moreover, it is not clear that we will typically satisfy step 5. That is, we may not usually find that an entity's having property  $\phi$  is explained by type-specific facts. We might find that having property  $\phi$  can be explained by the same sort of fact across entity types — e.g. members of two different entity types might share a property other than  $y$ , and *that* property could ground token instances of  $\phi$  in each case. Lastly, even when we do satisfy step 5, it does not follow that the property that grounds  $\phi$  in the case of entity  $x$  is  $y$ . That is, the justificatory process is not valid. The grounder could be some *other* property particular to members of entity type  $x$ . So I do not recommend the whole justificatory process here.

## Sum

I have examined the second-order views of three major figures in the grounding literature: Fine, Schaffer, and Audi. According to Fine, the constraints on grounding claims are intuition and explanatory power. To the extent that scientific considerations appear nowhere in Fine's second-order view, his view provides no guidance toward making grounding scientifically responsible. Schaffer's view might appear *prima facie* conducive to scientific responsibility, since one of the diagnostics of fundamentality claims mentions physics. However, neither that diagnostic, nor any of the others requires metaphysical theorizing to be done with an eye to science from the ground up, so his avowed methodology fails to secure scientific responsibility. Audi's suggested methodology does illuminate how grounding theories might be made scientifically responsible, since it shows how scientific explanation and evidence can be used in the justification of some grounding claims. Nevertheless, his justificatory process fails to be valid. So the grounding literature contains at least one example of a second-order view conducive to scientifically



responsible theories of grounding, but not one that I will recommend. Perhaps a more suitable method emerges in first-order practice.

## **I.2 Practice**

Since theory and practice sometimes diverge, and since not every philosopher ascends to the second order and makes her epistemology or her methodology explicit, I will now examine the first-order practices of metaphysicians working on grounding. Here I will organize my treatment around some of the central questions of the literature: whether reality is flat or hierarchical, whether there is a fundamental level, and what grounds what. Questions that I will not discuss include, *inter alia*, whether there is a single grounding relation and what the term *grounding* means. Regarding the latter, many grounding theorists (most notably Schaffer 2009) are silent. That is because they take *grounding* to be a primitive term that cannot be precisely analyzed and that can be elucidated only by providing examples, offering analogies, and “enfold[ing] the concept within a formalism” (Schaffer 2015, §1.1). Since few attempt to answer the question of the nature of grounding, the issue of making treatments of the question scientifically responsible does not arise. There is also a question of whether various uses of the notion of grounding equivocate on, or collapse into, other notions of dependence and whether the grounding concept uniquely performs indispensable explanatory work (see Wilson 2014). I take this question of unification to be answerable on conceptual grounds. This conceptual project can and ought to be done from the armchair. It is important theoretical work, but the impetus to go scientifically responsible is simply not applicable to this aspect of theory-building.

### I.2.1 Levels

Let us start with the question of whether reality is hierarchical or flat. The question is whether or not there are levels of reality that stand in asymmetric dependence relations to one another. For instance, there might be a mental level that depends on a biological level that depends on a mereologically atomic physical level. The question of whether grounding relations exist is relevant here, since, if the hierarchical view obtains, the asymmetrical ordering would be done by grounding relations. I will consider two reasons offered in support of a levelled picture of reality: 1) intuition favours it and 2) the structure of science supports it. Bennett takes the first route. She writes:

I have no knockdown argument against the claim that the world is flat. But every fiber of my being cries out in protest... *any* version of flatworldism will be radically revisionary... flatworldism is, to borrow a colorful word from a friend, ‘crazy pants’” (2011, 28).

Here, Bennett readily admits that her intuitions on this matter, however strong, do not constitute a *knockdown* argument. But does she think they constitute *an argument*? Does she think her intuitions provide *evidence* — defeasible evidence, but evidence nonetheless — of her favoured thesis? Bennett could simply be using intuition talk in the way that Dorr, in his response to Ladyman and Ross (2007), describes — i.e. as “a device for committing oneself to P while signaling that one is not going to provide any further arguments for this claim” (2010, np). That is, she may simply be signaling her rejection of flatworldism. This interpretation receives some support when Bennett says that her complaint against flatworldism is “a cousin of the incredulous stare” (2011, 28), because the incredulous stare is clearly not an argument. But then, why not simply *say*, plainly, ‘I will assume the falsity of flatworldism’? At any rate, we do not yet have an example of a scientifically responsible approach to the question — at best because this does not

constitute an attempt at justifying an answer to the question, at worst because it constitutes a free range attempt.

Do holistic considerations help Bennett here? Is her intuitive armchair claim vindicated by having a place in a theory that is scientifically responsible overall? It is not. Recall from Chapter 3 that holistic considerations can vindicate a metaphysical claim not directly supported by science only when: 1) the claim does not significantly constrain other components of the theory, or 2) the claim is itself constrained by other components of an overall scientifically responsible theory which themselves ultimately have scientific support. Bennett's claim about hierarchical structuring is too significant to treat in this holistic manner. How one answers it affects the course of inquiry — it constrains the questions one asks and the responses one can give to them. For instance, once one is committed to hierarchical structuring, one is then committed to a course of investigation regarding the structuring relations and relata, how to understand levels talk, questions of metaphysical and explanatory reduction, and so forth. If, on the other hand, one commits to flatworldism, one must account for levels of explanation, mereological relations, as well as emergent and supervenient phenomena, all without appeal to ontological levels. Moreover, the thesis of flatworldism is not constrained by components of an overall scientifically responsible metaphysics that ultimately have scientific support; it is constrained primarily by intuition. So holistic considerations do not vindicate the claim of flatworldism in this case.

By contrast, Schaffer (2003) argues by appeal to the structure of science. According to Schaffer, the hierarchical view of nature is “an empirical thesis based on the idea that the structure and discoveries of science reflects the structure of nature” (2003, 500). That is, we take as evidence the fact that our best sciences paint a structured picture of the world, and we draw the conclusion that the world is so structured. He explains: “grounding is a notion that is extremely natural in the

sciences, in considering the relation between levels. One need not be versed in an arcane metaphysics to think that the chemical depends on the physical” (2015, §4.4). In the picture of the world that emerges from science, the world is structured by these dependencies, and this is the metaphysician’s proper motivation and starting point for positing a hierarchical worldview structured by grounding relations. Here we have an approach to the question of whether reality is flat or hierarchical that is scientifically responsible, insofar as it looks to science and treats the scientific worldview as a source of evidence.

### **I.2.2 A Fundamental Level**

The grounding literature is also concerned with the question of whether the hierarchical structure of the world — the various levels and the asymmetric dependence relations that organize them — bottoms out in something absolutely fundamental. That is, whether the grounding chain stops in something ungrounded. This is a question analogous to the question of whether the causal nexus and the individual causal chains that compose it trace back to some uncaused cause. From Cameron (2008), we get a response to this question very similar to Bennett’s comments regarding levels. He reports:

I have identified an intuition, that... [t]here must be a metaphysical ground, a realm of ontologically independent objects which provide the ultimate ontological basis for all the ontologically dependent entities, and a realm of basic facts which provide the ultimate metaphysical grounding for all derivative facts. (2008, 8)

He attempts to argue for the intuition, but when his attempts fail, he says the following:

It is proving hard to argue for the intuition. Of course, it is an *intuition*; is not that reason enough to believe it? Yes, it is; I certainly feel the force of the intuition strongly, and I

think that this alone, given that I have seen no argument *against* the intuition, is sufficient to give me reason to believe that the intuition is true. We have to rest on intuition at *some* point, after all; is not here as good a place as any? Why should we even be trying to offer an argument for it? (2008, 11)

This is a much clearer case than the Bennett case above, in that Cameron clearly *does* see his intuitions as evidential. Cameron does, however, go on to provide an argument for fundamentality that he thinks *is* decisive, namely that positing a fundamental level provides “a common metaphysical explanation for every dependent entity” and therefore engenders a tidy picture of the world (2008, 12). Cameron’s claims here are free range; he makes no use of science and so his method fails to be scientifically responsible.

Moreover, any appeal to holistic considerations would fail in this case for precisely the reasons it failed in the Bennett case above. First, the fundamentality thesis constrains other aspects of the theory. It leads to an ontology of the fundamental and away from questions regarding infinite divisibility. Second, the fundamentality thesis is clearly not constrained by an overall scientifically responsible theory. It is constrained primarily by intuition and considerations of simplicity. So holistic considerations do not vindicate Cameron’s fundamentality claim.

In his treatment of the question of whether there is a fundamental level, Schaffer (2003) again does significantly better. Here, he does not take himself to be answering the question of whether there is *anything whatsoever* that is fundamental, but rather, whether there is a bottom-most, mereologically atomic level — a level consisting of units without proper parts (2003, 500). It is common to simply assume that there is a fundamental level, so understood, but Schaffer argues that science simply does not support the assumption. He says that when we ask whether there is a fundamental atomic structure in nature, we must “look to the structure of science for empirical

evidence” (2003, 500). He continues, “Thus I think that there is only one serious argument for atomism, which is that *science indicates atomism*—the trajectory of current physics points to a world constituted out of simple microphysical particles” (original emphasis, 2003, 500). The argument for atomism is:

- (1) There will be a complete microphysics,
- (2) The complete microphysics will postulate particles, and
- (3) These particles are the mereological atoms. (2003, 502)

In Schaffer’s view, this is the only serious argument *because* it is scientifically motivated. But he rejects it, also for scientifically motivated reasons. He rejects the argument because “the empirical evidence does not favor *any* of (1)-(3)” (2003, 503). In fact, he argues:

the history of science is a history of finding ever-deeper structure. We have gone from ‘the elements’ to ‘the atoms’ ... to the subatomic electrons, protons, and neutrons, to the zoo of ‘elementary particles’, to thinking that the hadrons are built out of quarks, and now we are sometimes promised that these entities are really strings, while some hypothesize that the quarks are built out of peons (in order to explain why quarks come in families). Should one not expect the future to be like the past. (2003, 503).

An induction on the trajectory of science would suggest that it will continue to posit further and further levels of physical reality. But ultimately, Schaffer concludes that with regard to the question of whether there is a fundamental level, “The empirical evidence to date is neutral as to which structure science is reflecting” (2003, 505). For that reason, he says, “one should withhold belief” (2003, 506). He considers whether the scientific evidence on offer supports or fails to support the fundamentality thesis and recommends the appropriate attitude; in so doing, his treatment of the question is scientifically responsible.

### I.2.3 What Grounds What

The question of what grounds what is the question of which sorts of entities or facts ground which other sorts of entities or facts, and of which token grounding relations obtain between particular entities or facts. In his treatment of the question, Audi claims:

Facts involving redness and loudness... never stand in grounding relations with one another. Nothing could be red in virtue of being loud, or loud in virtue of being red (and no range of intermediaries could ever link them together in a chain of grounding). These properties are simply too disparate. (2012, 108)

So, on his view, grounding facts are “importantly tied to the natures of properties” (2012, 108).

For instance, he says:

Compare maroonness. The fact that a thing is maroon is *bound* to ground its being red (assuming, for the moment, that a given thing’s redness is not identical with its maroonness). To label this relationship, let us say that facts are suited to stand in a relation of grounding only if their constituent properties are *essentially connected*. (2012, 108)

When an essential connection obtains between the constituent properties of the grounding relata, the type grounding relation is essential to the property type. That is, it is the nature of the one property to ground, or be grounded by, the other. So, if we find justification for positing an essential connection between properties, we are on our way to justifying a grounding claim. But what could justify belief about essential connection? Audi claims:

it *seems* to be of the essence of maroonness that its instances ground instances of redness (in the same particular). And likewise, it is of the essence of redness that an instance of it can be grounded in an instance of maroonness. (my emphasis, 2012, 109)

The added emphasis should signal that at least in this particular context, Audi appeals to intuition to settle the question of essential connection and therefore the question of what grounds what. So in this case, he does not follow the scientifically responsible justificatory procedure he articulated at the second order. But he could have done so! We have a good deal of scientific data pertaining to the nature of colour (see Akins and Hahn 2015 for a history of scientific development on the matter). One could easily appeal to scientific data about color properties to inform the discussion of essential connection in this case. But at least in this case, Audi justifies the essence claim on intuitive grounds, without any explicit consideration of science. So despite the scientific responsibility of Audi's second-order view, his first-order treatment of the what-grounds-what question fails to be scientifically responsible.

One particular version of the what-grounds-what question is the question of what is fundamental, since that question is about what ultimately grounds everything. Though Schaffer (2003) claims there is no scientific basis for positing *bottom-level, mereologically atomic* fundamentalia, he does not rule out other sorts of fundamentalia. He considers a number of arguments in favour of the claim that the whole cosmos, *qua* all-encompassing system, is the one and only fundamental object. This view is called *monism*. One of his arguments is from common sense — “the overall pull of intuitions across cultures and ages favors the monistic view” (2010, 49). Monism, he says, “properly understood as the claim that the cosmos is an integrated whole... best fits intuitions about priority” (2010, 50). The presence of an argument from intuition and common sense is unfortunate — but happily, Schaffer says, “I should hasten to add that I think this counts for little” (2010, 50). Instead, he focuses on articulating a scientifically motivated argument for monism. He says “advanced physics is field theoretic physics, and field theory has a natural monistic interpretation in terms of a spacetime bearing properties” (2009, 378). For



instance, he says: “general relativistic models are  $\langle M, g, T \rangle$  triples, where  $M$  is a four-dimensional continuously differentiable point manifold,  $g$  is a metric-field tensor, and  $t$  [*sic*] is a stress-energy tensor... The obvious ontology here is that of a spacetime manifold bearing fields” (2009, 378 fn 40). And further:

the cosmos forms one vast entangled system... Physically, one gets initial entanglement from the assumption that the world begins in one explosion (the Big Bang) in which everything interacts. This initial entanglement is then preserved thereafter on the assumption that the world evolves via Schrödinger’s equation. More precisely, the initial singularity is virtually certain... to produce universal entanglement, and the Schrödinger dynamics are virtually certain... to preserve it. In fact Schrödinger evolution tends to spread entanglements, so that even without initial entanglement, ‘eventually every particle in the universe must become entangled with every other’. (2010, 52)

These arguments present scientifically motivated reasons for accepting monism. He comments: “And so, given that quantum mechanics (or better, relativistic quantum field theory) represents our best current guide to the structure of reality, it seems that empirical inquiry now favors the holism of the monistic view” (2010, 55). His appeal to science is explicit. Since Schaffer takes the best current science and draws conclusions of philosophical interest from it, his argument for monism is scientifically responsible. Since, on his view, laws of metaphysics govern grounding relations (2010, 56), scientifically motivated token grounding claims could allow us to posit, in a scientifically responsible fashion, metaphysical laws about which types of things ground which types of things.

### I.2.4 Example-Based Method

Schaffer also makes a number of token grounding claims in a less scientifically responsible fashion. He proceeds via examples and counterexamples and considers what it is *plausible* or *natural* to say about those examples (see, for instance, his 2012, 126-129). In his view, this is partly how the concept of grounding must be communicated (2015, §1.1). In this respect, Schaffer is representative of a larger trend in the literature. Aside from Schaffer (2009), the example-based methodology is present in Fine (2012, 2011), Rosen (2010), and Audi (2012). Fine demonstrates the method most succinctly when he says: “There is an intuitive notion of one thing holding in virtue of another. Here are some examples...” (2012, 37). This example-based methodology is not scientifically responsible, because both the selection of the examples and the claims that one makes about them are guided primarily by intuition. We might say that our intuitions in this domain have some empirical weight, since they are generated by conceptual competence that we have achieved via experience. Even still, lacking any *scientific* basis for which examples we select and what we say about them, this example-based method is not scientifically responsible.

Do holistic considerations mitigate the problem? They don't. First, because the selection of paradigm examples significantly impacts the trajectory of the whole analysis, since any account of grounding must capture the paradigm cases and explain their features. Second, because the claims are not constrained by a scientifically responsible theory. They are constrained by something that is prior to a worked-out theory – an intuitive *grasp* of the concept of grounding, as Schaffer says (2009, 376). So holistic considerations do not vindicate claims produced by this example-based methodology.

Nevertheless, the work these examples do is necessary. Broadly speaking, the main components of theory-building include, at least:

- 1) formulating and clarifying *questions* of interest, and
- 2) formulating and justifying *responses* to those questions.

Then there may be a question of unifying our responses into a general, cohesive theoretical framework. At any rate, both 1) and 2) are crucial. These are not necessarily chronological, since, after doing some work to answer a question, we might get a better sense of the question. We can always return to the question, revise and reformulate it. But some treatment of 1) is a prerequisite for 2). The one is a groundwork for the other. Before we can answer a question, we must have at least some minimal understanding of what the question asks. The example-based methodology is a means of accomplishing the first component — a means of getting clear on questions of grounding by elucidating the concept of grounding. Again, this work is crucial, but the example-based methodology does not accomplish it in a scientifically responsible fashion.

But how else to proceed with this crucial work? Might we get our stock examples of grounding from science? Science might well furnish us with examples, but identifying which scientific dependencies are grounding relations will still be intuitive work — it will still require some antecedent concept of ground. So this crucial, preliminary conceptual work may be armchair work by necessity. And that is acceptable from the perspective of the advocate of scientific responsibility. It would be too strong a requirement to say that *all* aspects of scientifically responsible theorizing must proceed on the basis of science, especially if not all theoretical tasks are *apt* to be so performed. This aptness consideration distinguishes Schaffer's use of intuition from Bennett's and Cameron's. Bennett and Cameron appeal to intuition when science is demonstrably relevant. In contrast, Schaffer relies on intuition when his theoretical task is not clearly apt to be made scientifically responsible. So as a rule: scientifically responsible

metaphysical theories are supported on the whole by science, but science need not figure in at each and every stage of theory building.

Furthermore, that this particular component of theorizing does not appeal to science does not preclude its deliverances from being reinforced or undermined by science at a later stage in theorizing. A scientifically responsible treatment of a grounding question may reveal that the question is ill-formed. While conceptual work shapes the initial trajectory of an inquiry, engagement with science can ultimately correct that trajectory. It can show us when we have asked the wrong question or used the wrong concept. For instance, Darwinian advances in biology motivated us to change our typological conception of species, inherited from Plato, to a populational one – and this conceptual change makes a world of difference to our understanding of evolution (Mayr 2006).<sup>75</sup> In sum: a concept of ground is a prerequisite for a theory of ground, and it is not clear that science initially informs that concept. However, science can motivate conceptual revision later on. At any rate, theories of ground can nevertheless be scientifically responsible, so long as they accomplish most of their other tasks via engagement with science.

## **Sum**

In sum, I have examined treatments of a number of grounding questions, including the question of levels, of whether there is a fundamental level, and of what grounds what. In response to each question, we saw examples of first-order practices that fail to be scientifically responsible and of ones that succeed. Regarding the question of whether reality is flat or hierarchical, we saw Bennett report her strong intuition, without engaging science in her discussion of the question. In contrast, Schaffer argued that a structured view of reality emerges in science and it is *for that*

---

<sup>75</sup> I thank Derek Skillings for this example.

*reason* that the metaphysician should adopt a levelled picture. Regarding the question of whether there is a fundamental level, Cameron's arguments rely only on considerations of intuitive plausibility and simplicity, to the exclusion of any consideration of science. Again, Schaffer gives a scientifically responsible answer to the question when he considers whether available scientific evidence supports the fundamentality thesis, understood as the claim that there is a bottom-level, mereologically atomic level. Finally, regarding the question of what grounds what, we saw Audi posit a grounding claim on the basis of intuitive judgments about essential connections among properties. Schaffer again performs comparatively well, by arguing that monism is the fundamentality thesis most consistent with quantum mechanics. But Schaffer is also representative of a tendency in the literature toward an example-based methodology that relies on intuition. The evaluation of this first-order practice was less clear-cut. I concluded that, although it is not in itself scientifically responsible, the example-based methodology does not preclude a theory from being scientifically responsible on the whole, because not all theoretical tasks need to be accomplished via engagement with science. Moreover, it is possible that scientific engagement might correct errors made during the preliminary conceptual work. Generally, the first-order practices in the grounding literature are just as heterogeneous as one would expect. In some cases, the metaphysicians make dubious appeals to intuition in their treatment of the grounding questions. But among the proponents of grounding, Schaffer frequently demonstrates how grounding theories can be made scientifically responsible.

## **II. Scientifically Responsible Metaphysics of Ground**

On the basis of my examination of some of the first-order and second-order views espoused in the grounding literature, I will now synthesize some concrete recommendations for how to make

theories of ground scientifically responsible. The list comes in part from Audi's second-order justificatory process and Schaffer's first-order practice. In the formulation of theories of grounding, science can and ought to be used, *inter alia*:

- to help in the identification of putative grounding relata, by acquainting us with worldly objects, properties, states of affairs, and so on
- to show correlations among putative grounding relata
- to demonstrate the non-identity of relata
- to provide a stock of explanatory patterns
- to acquaint us with candidate essential properties and so help substantiate claims of essential connection and grounding
- to motivate and support the acceptance of metaphysical theses and the positing of metaphysical laws, e.g. about what is fundamental
- to motivate agnosticism about particular grounding theses where scientific support is lacking, e.g. about whether grounding chains descend infinitely
- to correct conceptual work performed independently of science and that turns out to be at odds with science

Some grounding claims will be subject to these forms of scientific support and some not. Value-laden grounding claims (like *the gods' love grounds the pious*) and grounding claims involving at least one abstract relatum (like *Socrates grounds Socrates' singleton*) may not be subject to direct scientific support. In these cases, we should adhere to something analogous to Ockham's Razor: we should not, beyond necessity, make grounding claims that lack direct scientific support. We should do so not for the purposes of simplicity, but because this type of claim is more *epistemically vulnerable* than scientifically supported ones and, as a matter of principle, we should minimize

epistemic vulnerability as much as possible. I have also claimed that holism can mitigate the epistemic risk of making grounding claims that lack direct scientific support, but only when such claims do not significantly constrain other components of the theory or are themselves constrained by an overall scientifically responsible theory.

In addition to these recommendations that I have drawn out of the literature, I also wish to gesture toward a scientifically responsible method for positing and testing *token* grounding claims. The method is akin to causal intervention tests on putatively causal correlates. Causal relations can be posited and tested on the basis of interventions on putative causes. This method for making and testing causal claims is empirical insofar as it tracks observable phenomena, or the observable consequences of unobservable phenomena. The evidence is scientific when it originates in the context of scientific experimentation. Grounding is a close cousin of causation, since both are asymmetric dependence relations that back explanation, so the method for justifying claims about the one may apply to the other. That is, grounding claims can be motivated and tested by appeal to empirical, indeed *scientific*, counterfactual dependency tests. When putative grounding relations are observable or have observable correlates, we can intervene on the putative grounder and check for a result in the putatively grounded entity (or their correlates). Here we have a scientific method for the formulation of token grounding claims. Type grounding claims may be inferred after grounding has been shown in sufficiently many relevantly similar cases.

Schaffer (2015) has recently suggested something similar. He argues that causation and grounding are alike, since:

- both are generative relations;
- both are partial orders, admitting a type/token distinction, a component/net distinction, an incomplete/complete/total distinction, and screening-off relations; and

- both are backed by non-accidental generalizations, delimit a specific form of necessity, are supportive of and diagnosable by counterfactuals, and can back explanation. (2015, §3.1)

In addition, he argues that both grounding and causation are “best formalized via structural equation models” (2015, §1.4). He points out that “within a restricted simple range of cases, there is wide agreement that counterfactual dependence tracks token causation” (2015, §2.3). He then argues that “there is a straightforward and informative parallel working test of token grounding to be had, in terms of counterfactual covariation: wiggle the ground, and the grounded wiggles” (Schaffer 2015, §3.2). To be scientifically respectable, the ‘wiggling’ would have to be an experimental activity, not merely an imaginative one.

However, the suggestion has limitations and its details need to be worked out. First, the test will only apply to select cases. That is because many putative grounders, such as the gods’ love and the pious, are not susceptible to causal intervention. Relatedly, it is not clear that the wiggling of the grounded entity will always have empirically tractable consequences. For instance, here’s a grounding claim that I take it Schaffer would accept: Santa Claus (a non-existent, non-concrete thing) is grounded in the whole cosmos.<sup>76</sup> We shouldn’t expect any wiggling of the cosmos to make Santa Claus wiggle. Further, we would have to spell out what the interventions would consist in and the sorts of effects we would be looking for. After identifying putative relata, observing correlations among them, and establishing non-identity, all in a scientific manner, we might remove (destroy or alter) the putative grounder and check whether the other one remains. But it is difficult to fill in the details while remaining neutral on the question of the relata of the grounding relation. So different theorists will have to fill in the details appropriately. The last

---

<sup>76</sup>Schaffer takes fictional beings to exist derivatively. See his discussion of God (2009, 359).



difficulty is that it is not clear how to differentiate token causal relations from token grounding relations. If the test is a causal intervention designed to show counterfactual dependence, one would think the results would show causation rather than grounding. But perhaps the difference will be in the relata – the causal relata will be events, the grounding relata something else, probably facts or properties or objects.

## II.1 Hofweber's Challenge

Before concluding, I wish to discharge one of Hofweber's challenges to grounding, which *prima facie* threatens the prospects of a scientifically responsible grounding literature. Hofweber thinks that most examples of metaphysical priority are *really* examples of something else: logical or conceptual priority or even causal priority (2009, 269-271). I remain neutral on this point, though I think it is important and worrying. But the point I will presently concern myself with is the following. According to Hofweber, if we posit some special kind of metaphysical priority, then “[w]e have autonomy from the facts” (2009, 273). What we say about metaphysical priority is not constrained by empirical facts *whatsoever*. He explains: “Even though some who hold on to metaphysical priority, like Schaffer, think that science tracks what is prior in this sense, this isn't a requirement at all” (2009, 273). That is, metaphysical priority, if there is such a thing, need not track scientific priority. He continues: “Why not think that what science tracks is merely what is scientifically prior, which might or might not coincide with what is metaphysically prior?” (2009, 273). On such a view, philosophy “investigates what the world is REALLY like” while the sciences “only find out what the world is like” (2009, 273). In other words, once you countenance metaphysical priority, nothing restricts the metaphysician to empirical facts or to scientific priority. She can acknowledge the facts of scientific priority and then go on to posit a *deeper*,

metaphysically prior layer of structure, and claim *that's* what is *ultimately* fundamental, or *ultimately* real. The door is open to posit virtually anything she likes. For instance, Hofweber muses, “I can’t wait for the first metaphysician to come out and defend... *priority aquaism*: everything is ultimately water. Water is the most fundamental of all things”, which is “[n]ot to be confused with *aquaism*: the view that everything is water” (2009, 273). Such a view might have some virtues: “It nicely goes with a process metaphysics. It supports our intuitive judgment that water is an especially important liquid... Maybe it even gives rise to the final explanation of why time flows” (2009, 273). So, the criticism goes, just look what can happen if we allow this notion of metaphysical priority! If the criticism is apt, then the grounding literature cannot be made scientifically responsible, because when it posits metaphysical priority, it suggests that “philosophy is really queen of the sciences after all” (2009, 273) and is ultimately autonomous from science.

I respond that Hofweber has misdescribed the typical view of metaphysical priority. On a charitable reading of the grounding literature, metaphysicians are not positing some *sui generis* structuring relation *over and above* scientific priority. Rather, they claim there is *one* structure in the world — *one* set of priority relations. Call them metaphysical priority relations. Some of them will be scientifically discoverable and some not. The metaphysician does not posit an additional or separate layer of structure. Hofweber is right to be concerned about views like *priority aquaism*, but what the example shows is not that positing metaphysical priority leads inescapably to absurdity and theoretical anarchy; it shows, simply, the need for scientific responsibility. The example demonstrates precisely what I argued for in Chapter 3: the sheer epistemic recklessness of relying solely on free range theoretical constraints — such as coherence with antecedently accepted commitments, intuitive plausibility, and explanatory power — in the formulation and

justification of substantive metaphysical theses. It demonstrates the need for a robust, justificatory constraint on our theorizing — namely, a scientific constraint.

### **III. Conclusion, Case II**

I have performed a case study of the grounding literature, with an eye to the prospect of making it scientifically responsible. I did a piecemeal analysis of various influential works in the literature, starting with explicit second-order commitments (epistemology and methodology) and proceeding to first-order practices (metaphysics). I found the second-order views and first-order practices of metaphysicians in this subfield to be, unsurprisingly, heterogeneous.

At the second order, I found that Fine's avowed epistemology excludes considerations of science. With its mention of physics, Schaffer's second-order view *prima facie* points the way toward making grounding theories scientifically responsible, but it turns out to require mere consistency with science, not conscientious engagement with it. Audi's justificatory process is the most promising, since it encourages the use of scientific evidence and explanation in the justification of some grounding claims, but it fails in other respects.

At the first order, Bennett, Cameron and Audi report intuitions, without relevant scientific evidence. In contrast, Schaffer demonstrates what scientifically responsible grounding theories can look like, by using scientific evidence to motivate belief in levels and monism, as well as neutrality toward mereological atomism. However, he is also representative of a trend in the literature toward an example-based methodology that is not itself scientifically responsible. But I concluded that preliminary conceptual work might be a properly armchair task and might not undermine the epistemic credentials of the overall theory. Moreover, the example-based methodology does not preclude science correcting the concepts or paradigm cases at a later stage in the theorizing.

From my examination of the literature, I synthesized a number of general recommendations for how theorists of ground might fruitfully engage science. I also gestured toward a method for positing and testing grounding claims involving intervention on putative grounding relata — though the details of this method remain to be worked out. Finally, I considered Hofweber’s suggestion that, in positing metaphysical priority, grounding theorists proclaim autonomy from science. In response, I argued that metaphysical priority is not meant to be a *sui generis* form of priority over and above scientific priority. Moreover, what Hofweber points to is not the *impossibility* of making grounding theories answerable to science, but the *need* to do so.

## CONCLUSION

I wrote this dissertation out of broad sympathy with those who have sounded a clarion call for the naturalization of metaphysics, but with the conviction that two crucial things are missing from the relevant literature – namely, an adequate account of: 1) what makes a metaphysical theory naturalized or not, and 2) why the naturalization of metaphysics is a good thing. The broadest aim of this dissertation was to account for those two things.

Regarding 1), I claimed that metaphysics' being naturalized is a function of its engaging conscientiously with the theories and practices of the current sciences, where *science* is to be understood ostensibly, *conscientious* engagement rules out independent theorizing that looks to science only as an afterthought, and *engagement* means potentially many things, including unifying scientific theses, integrating bits of science into metaphysical theories, drawing philosophical conclusions from scientific evidence or practice, interpreting scientific data, correcting scientists' interpretations, and revising metaphysical claims or questions in light of new scientific evidence. The more metaphysics engages science in these ways, the more naturalized it is.

Regarding 2), I argued that robust theoretical constraint and epistemic warrant conduce to a number of different putative epistemic goods, including greater statistical likeliness, agreement, falsity avoidance, methodological expediency and systematicity, truth-conduciveness, relevance, and rationality. I concluded that for epistemic purposes, we should robustly constrain our theories and seek adequate epistemic warrant for them. I then argued that non-naturalized or free range metaphysics fails to do this. Naturalizing it or making it scientifically responsible resolves the problem, since scientifically responsible metaphysics inherits some of the constraint and epistemic

warrant that science enjoys. So naturalizing metaphysics is a good thing because it makes metaphysics epistemically adequate when it isn't otherwise.

Metaphysics is and will continue to be pursued via a number of methods, both free range and scientifically responsible. Both have their uses. However, I take myself to have shown that where we face a choice between free range and scientifically responsible metaphysics *and* where we care to carve out a view that we are epistemically justified in accepting, we must take the more difficult road and look to science as a source of evidence and as a constraint on the kinds of claims we make. I am therefore greatly encouraged by the recent surge of interest and activity in what's now called scientific metaphysics or the metaphysics of science. If metaphysics is to be more than a toolbox, but a substantive contributor to the human quest to understand and describe the world, that's where its contributions will be.

## BIBLIOGRAPHY

- Adler, Jonathan. 2002. *Belief's Own Ethics*. Cambridge, MA: MIT Press.
- Ahlstrom-Vij, Kristoffer. 2013. "In Defense of Veritistic Value Monism." *Pacific Philosophical Quarterly* 94 (1): 19–40.
- , and Jeffrey Dunn, eds. forthcoming. *Epistemic Consequentialism*. Oxford: Oxford UP.
- , and Jeffrey Dunn. 2014. "A Defence of Epistemic Consequentialism." *Philosophical Quarterly* 64 (257): 541–551.
- Akins, Kathleen, and Martin Hahn. 2015. "Colour Perception." In *The Oxford Handbook of the Philosophy of Perception*, edited by Mohan Matthen, 422–440. Oxford: Oxford UP.
- Albert, David. 2012. "On the Origin of Everything." *New York Times*.  
<http://www.nytimes.com/2012/03/25/books/review/a-universe-from-nothing-by-lawrence-m-krauss.html>.
- . 2015. *After Physics*. Cambridge, MA: Harvard UP.
- Alicke, Mark, David Rose, and Dori Bloom. 2011. "Causation, Norm Violation, and Culpable Control." *Journal of Philosophy* 108 (12): 670–696.
- Allori, Valia. 2015. "Primitive Ontology in a Nutshell." *International Journal of Quantum Foundations* 1 (2): 107–122.
- Alston, William. 1985. "Concepts of Epistemic Justification." *The Monist* 68 (1): 57–89.
- . 1986. "Epistemic Circularity." *Philosophy and Phenomenological Research* 47 (1): 1–30.
- . 1988. "The Deontological Conception of Epistemic Justification." *Philosophical Perspectives* 2: 257–299.
- . 1989. *Epistemic Justification*. Ithaca, NY: Cornell University Press.

- Andow, James. 2016. "Reliable But Not Home Free? What Framing Effects Mean for Moral Intuitions." *Philosophical Psychology* 6: 1–8.
- Anselm. 1078. *Proslogion*.
- Armstrong, David. 1973. *Belief, Truth and Knowledge*. Cambridge: Cambridge UP.
- Audi, Paul. 2012. "A Clarification and Defense of the Notion of Grounding." In *Metaphysical Grounding: Understanding the Structure of Reality*, edited by Fabrice Correia and Benjamin Schnieder, 101–121. Cambridge: Cambridge UP.
- Audi, Robert. 1988. "Justification, Truth, and Reliability." *Philosophy and Phenomenological Research* 49 (1): 1–29.
- . 2001. "Doxastic Voluntarism and the Ethics of Belief." In *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*, edited by Matthias Steup, 93–111. Oxford: Oxford UP.
- Axtell, Guy. 1997. "Recent Work on Virtue Epistemology." *American Philosophical Quarterly* 34 (1): 1–26.
- , ed. 2000. *Knowledge, Belief, and Character: Readings in Virtue Epistemology*. Lanham, MD: Rowman and Littlefield.
- , and J. Adam Carter. 2008. "Just the Right Thickness: A Defense of Second-Wave Virtue Epistemology." *Philosophical Papers* 37 (3): 413–434.
- Baehr, Jason. 2009. "Is There a Value Problem?" In *Epistemic Value*, edited by Adrian Haddock, Alan Millar, and Duncan Pritchard, 42–59. Oxford: Oxford UP.
- Barker, Matthew. 2010. "Species Intrinsicism." *Philosophy of Science* 77: 73–91.



- Bealer, George. 2002. "Modal Epistemology and the Rationalist Renaissance." In *Conceivability and Possibility*, edited by Tamar Gendler and John Hawthorne, 71–126. Oxford: Oxford UP.
- Bennett, Karen. 2009. "Composition, Colocation, and Metaontology." In *Metametaphysics: New Essays on the Foundations of Ontology*, edited by David Chalmers, David Manley, and Ryan Wasserman, 38–76. Oxford: Oxford UP.
- . 2011. "By Our Bootstraps." *Philosophical Perspectives* 25 (Metaphysics): 27–41.
- . 2016. "There Is No Special Problem with Metaphysics." *Philosophical Studies* 173: 21–37.
- Berger, James. 1985. *Statistical Decision Theory and Bayesian Analysis*. New York: Springer-Verlag.
- Berker, Selim. 2013. "The Rejection of Epistemic Consequentialism." *Philosophical Issues* 23 (1): 363–387.
- Betz, Gregor. 2013. "Degrees of Justification, Bayes' Rule, and Rationality." In *Bayesian Argumentation: The Practical Side of Probability*, edited by Frank Zenker, 135–146. New York: Springer.
- Bird, Alexander. 2007. "Justified Judging." *Philosophy and Phenomenological Research* 74 (1): 81–110.
- Bishop, Robert. 2002. "Chaos, Indeterminism, and Free Will." In *The Oxford Handbook of Free Will*, edited by Robert Kane, 111–126. New York: Oxford UP.
- Boghossian, Paul. 2006. *Fear of Knowledge: Against Relativism and Constructivism*. Oxford: Oxford UP.
- BonJour, Laurence. 1985. *The Structure of Empirical Knowledge*. Cambridge, MA: Harvard UP.

- Bouzat, Juan. 2014. "Darwin's Diagram of Divergence of Taxa as a Causal Model for the Origin of Species." *Quarterly Review of Biology* 89 (1): 21–38.
- Box, George E.P., and George C. Tiao. 1973. *Bayesian Inference and Statistical Analysis*. New York: John Wiley & Sons.
- Briesen, Jochen. 2016. "Epistemic Consequentialism: Its Relation to Ethical Consequentialism and the Truth-Indication Principle." In *Epistemic Reasons, Norms and Goals*, edited by Pedro Schmechtig and Martin Grajner, 277–306. Berlin: De Gruyter.
- Brogaard, Berit. 2008. "The Trivial Argument for Epistemic Value Pluralism, Or How I Learned to Stop Caring About Truth." In *Epistemic Value*, edited by Adrian Haddock, Alan Millar, and Duncan Pritchard, 284–305. Oxford: Oxford UP.
- Bruno, Michael, and Shaun Nichols. 2010. "Intuitions about Personal Identity: An Empirical Study." *Philosophical Psychology* 23 (3): 293–312.
- Bueno, Otávio, and Scott Shalkowski. 2014. "Modalism and Theoretical Virtues: Toward an Epistemology of Modality." *Philosophical Studies* 172 (3): 671–689.
- Callender, Craig. 2011. "Philosophy of Science and Metaphysics." In *The Continuum Companion to The Philosophy of Science*, edited by Steven French and Juha Saatsi, 33–54. New York: Continuum.
- Calosi, Claudio, and Matteo Morganti. 2016. "Humean Supervenience, Composition as Identity and Quantum Wholes." *Erkenntnis* 81 (6): 1173–1194.
- Cameron, Ross. 2008. "Turtles All the Way Down: Regress, Priority and Fundamentality." *Philosophical Quarterly* 58 (230): 1–14.
- Carnap, Rudolf. 1950. "Empiricism, Semantics, and Ontology." *Revue Internationale de Philosophie* 4 (11): 20–40.

- Chalmers, David. 1996. *The Conscious Mind: In Search of a Fundamental Theory*. Oxford: Oxford UP.
- . 2002. “Does Conceivability Entail Possibility?” In *Conceivability and Possibility*, edited by Tamar Gendler and John Hawthorne, 145–200. Oxford: Oxford UP.
- . 2009. “Ontological Anti-Realism.” In *Metametaphysics: New Essays on the Foundations of Ontology*, edited by David Chalmers, David Manley, and Ryan Wasserman, 77–129. Oxford: Oxford UP.
- , David Manley, and Ryan Wasserman, eds. 2009. *Metametaphysics: New Essays on the Foundations of Ontology*. Oxford: Oxford UP.
- Chang, Hasok. 2012. *Is Water H<sub>2</sub>O? Evidence, Realism and Pluralism*. New York: Springer.
- Chennu, Srivas, Valdas Noreika, David Gueorguiev, Yury Shtyrov, Tristan Bekinschtein, and Richard Henson. 2016. “Silent Expectations: Dynamic Causal Modeling of Cortical Prediction and Attention to Sounds That Weren’t.” *The Journal of Neuroscience* 36 (32): 8305–8316.
- Cherniak, Christopher. 1986. *Minimal Rationality*. Cambridge, MA: MIT Press.
- Chisholm, Roderick. 1980. “A Version of Foundationalism.” *Midwest Studies in Philosophy* 5 (1): 543–564.
- . 1989. *Theory of Knowledge*. 3rd ed. Englewood Cliffs, NJ: Prentice-Hall.
- Chrisman, Matthew. 2008. “Ought to Believe.” *Journal of Philosophy* 105 (7): 346–370.
- Chudnoff, Elijah. 2014. *Intuition*. Oxford: Oxford UP.
- Coates, Allen. 2009. “Explaining the Value of Truth.” *American Philosophical Quarterly* 46 (2): 105–115.
- Cohen, L. Jonathan. 1992. *An Essay on Belief and Acceptance*. Oxford: Clarendon Press.

- Colombo, Matteo. 2016. "Experimental Philosophy of Explanation Rising: The Case for a Plurality of Concepts of Explanation." *Cognitive Science* 40 (4): 1–15.
- Conee, Earl, and Richard Feldman. 2001. "Internalism Defended." In *Epistemology: Internalism and Externalism*, edited by Hilary Kornblith, 231–260. Malden, MA: Blackwell.
- , and Richard Feldman. 2004. *Evidentialism*. Oxford: Oxford UP.
- , and Ted Sider. 2005. *Riddles of Existence: A Guided Tour of Metaphysics*. Oxford: Oxford UP.
- Costa, Newton, Olimpia Lombardi, and Mariano Lastiri. 2013. "A Modal Ontology of Properties for Quantum Mechanics." *Synthese* 190 (17): 3671–3693.
- Crane, Tim. 2013. *The Objects of Thought*. Oxford: Oxford UP.
- Cruz, Joe, and John Pollock. 2004. "The Chimerical Appeal of Epistemic Externalism." In *The Externalist Challenge*, edited by Richard Schantz, 125–142. Berlin: De Gruyter.
- Daly, Chris. 2012. "Scepticism About Grounding." In *Metaphysical Grounding: Understanding the Structure of Reality*, edited by Fabrice Correia and Benjamin Schnieder, 81–100. Cambridge: Cambridge UP.
- David, Marian. 2013. "Truth as the Primary Epistemic Goal: A Working Hypothesis." In *Contemporary Debates in Epistemology*, edited by Matthias Steup, John Turri, and Ernest Sosa, Second Edi, 363–377. Malden, MA: Wiley Blackwell.
- Dawid, Richard. 2013. "Theory Assessment and Final Theory Claim in String Theory." *Foundations of Physics* 43 (81): 81–100.
- DePaul, Michael. 2001. "Value Monism in Epistemology." In *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*, edited by Matthias Steup, 170–183. Oxford: Oxford UP.

- Descartes, Rene. 1637. *Discourse on Method*.
- Devitt, Michael. 1996. *Coming to Our Senses: A Naturalistic Program for Semantic Localism*.  
Cambridge: Cambridge UP.
- . 1997. *Realism and Truth*. 2nd Ed. Princeton: Princeton UP.
- . 1998. “Naturalism and the A Priori.” *Philosophical Studies* 92 (1/2): 45–65.
- . 2008. “Resurrecting Biological Essentialism.” *Philosophy of Science* 75 (3): 344–382.
- . 2010. *Putting Metaphysics First: Essays on Metaphysics and Epistemology*. Oxford:  
Oxford UP.
- . 2011. “No Place for the A Priori.” In *What Place for the A Priori?*, edited by Michael  
Shaffer and Michael Veber, 9–32. Chicago: Open Court.
- . 2012. “Whither Experimental Semantics?” *Theoria* 27 (1): 5–36.
- . 2014. “We Don’t Learn About the World by Examining Concepts: A Response to Carrie  
Jenkins.” In *Current Controversies in Epistemology*, edited by Ram Neta, 23–34. New  
York: Routledge.
- . 2015. “Relying on Intuitions: Where Cappelen and Deutsch Go Wrong.” *Inquiry* 58 (7–  
8): 669–699.
- Dorato, Mauro. 2015. “Events and the Ontology of Quantum Mechanics.” *Topoi* 34 (2): 369–  
378.
- , and Matteo Morganti. 2013. “Grades of Individuality: A Pluralistic View of Identity in  
Quantum Mechanics and in the Sciences.” *Philosophical Studies* 163 (3): 591–610.
- Dorr, Cian. 2010. “Review of Every Thing Must Go: Metaphysics Naturalized.” *Notre Dame  
Philosophical Review*. <https://ndpr.nd.edu/news/24377-every-thing-must-go-metaphysics-naturalized/>.

- Duhem, Pierre. 1914. *La Théorie Physique: Son Objet et Sa Structure*. Paris: Marcel Riviera & Cie.
- Dunn, Jeffrey. 2016. "Epistemic Consequentialism." *The Internet Encyclopedia of Philosophy*. Accessed December 2. <http://www.iep.utm.edu/epis-con/>.
- Dupré, John. 1993. *The Disorder of Things: Metaphysical Foundations of the Disunity of Science*. New York: Harvard UP.
- Dyke, Heather, and James Maclaurin. 2013. "What Shall We Do with Analytic Metaphysics? A Response to McLeod and Parsons." *Australasian Journal of Philosophy* 91 (1): 179–182.
- Ellis, Brian. 2001. *Scientific Essentialism*. Cambridge: Cambridge UP.
- . 2002. *The Philosophy of Nature: A Guide to the New Essentialism*. Montreal: McGill-Queen's UP.
- Ellis, George, and Joe Silk. 2014. "Scientific Method: Defend the Integrity of Physics." *Nature* 516: 321–323. [www.nature.com/news/scientific-method-defend-the-integrity-of-physics-1.16535](http://www.nature.com/news/scientific-method-defend-the-integrity-of-physics-1.16535).
- Engel, Pascal. 2007. "Belief and Normativity." *Disputatio* 2 (23): 179–203.
- Ereshefsky, Marc. 1991. "Species, Higher Taxa, and the Units of Evolution." *Philosophy of Science* 58 (1): 84–101.
- . 1994. "Some Problems with the Linnaean Hierarchy." *Philosophy of Science* 61 (2): 186–205.
- . 1998. "Species Pluralism and Anti-Realism." *Philosophy of Science* 65: 103–120.
- . 2001. *The Poverty of the Linnaean Hierarchy: A Philosophical Study of Biological Taxonomy*. Cambridge: Cambridge UP.

- . 2010a. “Microbiology and the Species Problem.” *Biology and Philosophy* 25 (4): 553–568.
- . 2010b. “What’s Wrong with the New Biological Essentialism.” *Philosophy of Science* 77 (5): 674–685.
- , and Makmiller Pedroso. 2013. “Biological Individuality: The Case of Biofilms.” *Biology and Philosophy* 28 (2): 331–349.
- Esfeld, Michael. 2004. “Quantum Entanglement and a Metaphysics of Relations.” *Studies in History and Philosophy of Science Part B* 35 (4): 601–617.
- , Dustin Lazarovici, Mario Hubert, and Detlef Dürr. 2014. “The Ontology of Bohmian Mechanics.” *British Journal for the Philosophy of Science* 64 (4): 773–796.
- Feldman, Richard. 1988. “Epistemic Obligations.” *Philosophical Perspectives* 2: 235–256.
- . 2001. “Voluntary Belief and Epistemic Evaluation.” In *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*, edited by Matthias Steup, 77–92. Oxford: Oxford UP.
- . 2002. “Epistemological Duties.” In *The Oxford Handbook of Epistemology*, 362–384.
- Fine, Kit. 2001. “The Question of Realism.” *Philosophers’ Imprint* 1 (1): 1–30.
- . 2012. “Guide to Ground.” In *Metaphysical Grounding*, edited by Fabrice Correia and Benjamin Schnieder, 37–80. Cambridge: Cambridge UP.
- Fischer, Bob, and Felipe Leon, eds. 2017. *Modal Epistemology After Rationalism*. Dordrecht: Springer.
- Fodor, Jerry. 1998. *Concepts: Where Cognitive Science Went Wrong*. Oxford: Oxford UP.
- Foley, Richard. 1987. *The Theory of Epistemic Rationality*. Cambridge, MA: Harvard UP.
- . 1993. *Working Without a Net*. Oxford: Oxford UP.

- Forster, Malcolm, and Elliott Sober. 1994. "How to Tell When Simpler, More Unified, or Less Ad Hoc Theories Will Provide More Accurate Predictions." *British Journal for the Philosophy of Science* 45: 1–35.
- Fox, John. 1988. "The Ethnomethodology of Science." In *Realism and Relativism in Science*, edited by Robert Nola, 59–80. Dordrecht: Kluwer.
- French, Steven. 1998. "On the Withering Away of Physical Objects." In *Interpreting Bodies: Classical and Quantum Objects in Modern Physics*, edited by Elena Castellani, 93–113. Princeton: Princeton UP.
- . 2014. *The Structure of the World: Metaphysics and Representation*. Oxford: Oxford UP.
- , and Kerry McKenzie. 2012. "Thinking Outside the Toolbox: Towards a More Productive Engagement Between Metaphysics and Philosophy of Physics." *European Journal of Analytic Philosophy* 8 (1): 42–59.
- Frise, Matthew. 2017. "No Need to Know." *Philosophical Studies* 174 (2): 391–401.
- Geirsson, Heimir. 2005. "Conceivability and Defeasible Modal Justification." *Philosophical Studies* 122 (3): 279–304.
- Ghiselin, Michael. 1974. "A Radical Solution to the Species Problem." *Systematic Zoology* 47: 350–383.
- Gibbard, Allan. 2012. *Meaning and Normativity*. Oxford: Oxford UP.
- Giere, Ronald. 1988. *Explaining Science: A Cognitive Approach*. Chicago: University of Chicago Press.
- . 1989. "Scientific Rationality as Instrumental Rationality." *Studies in History and Philosophy of Science* 20: 377–384.
- Ginet, Carl. 1975. *Knowledge, Perception, and Memory*. Dordrecht: D. Reidel Pub. Co.



- . 2001. “Deciding to Believe.” In *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*, edited by Matthias Steup, 63–76. Oxford: Oxford UP.
- Giret, Nicolas, Joergen Kornfeld, Ganguli Surya, and Richard Hahnloser. 2014. “Evidence for a Causal Inverse Model in an Avian Cortico-Basal Ganglia Circuit.” *Proceedings of the National Academy of Sciences of the United States of America* 111 (16): 6063–6068.
- Godfrey-Smith, Peter. 2012. “Metaphysics and the Philosophical Imagination.” *Philosophical Studies* 160 (1): 97–113.
- . 2013. *On the Relation Between Philosophy and Science*.  
[http://petergodfreysmith.com/wp-content/uploads/2013/06/PhilosophyScience\\_PGS\\_2013\\_C.pdf](http://petergodfreysmith.com/wp-content/uploads/2013/06/PhilosophyScience_PGS_2013_C.pdf).
- . 2014. *Philosophy of Biology*. Princeton, NJ: Princeton UP.
- Goldman, Alvin. 1967. “A Causal Theory of Knowing.” *Journal of Philosophy* 64 (12): 357–372.
- . 1979. “What Is Justified Belief?” In *Justification and Knowledge*, edited by George Pappas, 1–25. Boston: D. Reidel Pub. Co.
- . 1999. *Knowledge in a Social World*. Oxford: Oxford UP.
- . 2002. “The Unity of the Epistemic Virtues.” In *Pathways to Knowledge*, 52–70. Oxford: Oxford UP.
- . 2007. “Philosophical Intuitions: Their Target, Their Source, and Their Epistemic Status.” *Grazer Philosophische Studien* 74: 1–26.
- . 2015. “Reliabilism, Veritism, and Epistemic Consequentialism.” *Episteme* 12 (2): 131–143.

- Graham, Peter. 2011. "Does Justification Aim at Truth?" *Canadian Journal of Philosophy* 41 (1): 51–72.
- Grant, Robert. 2011. "Naturalizing the Metaphysics of Species: A Perspective on the Species Problem." *South African Journal of Philosophy* 30 (1): 63–69.
- Greco, John. 2010. *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity*. Cambridge: Cambridge UP.
- Grimm, Stephen. 2009. "Epistemic Normativity." In *Epistemic Value*, edited by Adrian Haddock, Alan Millar, and Duncan Pritchard, 243–264. Oxford: Oxford UP.
- Gross, Paul, and Norman Levitt. 1998. *Higher Superstition: The Academic Left and Its Quarrels with Science*. Baltimore: John Hopkins UP.
- Haack, Susan. 2003. *Defending Science - Within Reason: Between Scientism and Cynicism*. Amherst: Prometheus Books.
- Hale, Bob. 2002. "Knowledge of Possibility and of Necessity." *Proceedings of the Aristotelian Society* 103 (1): 1–20.
- Hanrahan, Rebecca Roman. 2009. "Consciousness and Modal Empiricism." *Philosophia* 37: 281–306.
- Hazlett, Allan. 2013. *A Luxury of the Understanding: On the Value of True Belief*. Oxford: Oxford UP.
- Henderson, Leah, Noah Goodman, Joshua Tenenbaum, and James Woodward. 2010. "The Structure and Dynamics of Scientific Theories: A Hierarchical Bayesian Perspective." *Philosophy of Science* 77 (2): 172–200.
- Henning, Tim, and David Schweikard, eds. 2013. *Knowledge, Virtue, and Action: Putting Epistemic Virtues to Work*. New York: Routledge.

- Hieronymi, Pamela. 2008. "Responsibility for Believing." *Synthese* 161 (3): 357–373.
- Hitchcock, Christopher, and Joshua Knobe. 2009. "Cause and Norm." *Journal of Philosophy* 106 (11): 587–612.
- Hodgson, David. 2002. "Quantum Physics, Consciousness, and Free Will." In *The Oxford Handbook of Free Will*, edited by Robert Kane, 85–110. New York: Oxford UP.
- Hofweber, Thomas. 2009. "Ambitious, Yet Modest, Metaphysics." In *Metametaphysics: New Essays on the Foundations of Ontology*, edited by David Chalmers, David Manley, and Ryan Wasserman, 260–289. Oxford: Oxford UP.
- Huemer, Michael. 2007. "Moore's Paradox and the Norm of Belief." In *Themes from G.E. Moore*, 142–157. Oxford: Oxford UP.
- Hull, David. 1978. "A Matter of Individuality." *Philosophy of Science* 45: 335–360.
- Jäger, Christoph. 2004. "Epistemic Deontology, Doxastic Voluntarism, and the Principle of Alternate Possibilities." In *Knowledge and Belief*, edited by Winfried Löffler and Paul Weingartner, 217–227. Wien: ÖBV.
- Jeffreys, Harold. 1939. *The Theory of Probability*. Oxford: Oxford UP.
- Jeffreys, William, and James Berger. 1992. "Ockham's Razor and Bayesian Analysis." *American Scientist* 80 (1): 64–72.
- Jenkins, Carrie. 2010. "Concepts, Experience and Modal Knowledge." *Philosophical Perspectives* 24: 255–279.
- Kane, Robert. 1996. *The Significance of Free Will*. Oxford: Oxford UP.
- Kant, Immanuel. 1785/2001. *Prolegomena to Any Future Metaphysics*. Edited by James Ellington. 2nd edition. Indianapolis: Hackett.

- Kellert, Stephen, Helen Longino, and Ken Waters, eds. 2006. *Scientific Pluralism*. Minnesota: University of Minnesota Press.
- Kelly, Kevin. 2004. "Justification as Truth-Finding Efficiency: How Ockham's Razor Works." *Minds and Machines* 14: 485–505.
- Kistler, Max. 2010. "Mechanisms and Downward Causation." *Philosophical Psychology* 22 (5): 595–609.
- Kitcher, Philip. 1981. "Explanatory Unification." *Philosophy of Science* 48 (4): 507–531.
- . 1982. *Abusing Science: The Case Against Creationism*. Cambridge, MA: MIT Press.
- . 1989. "Explanatory Unification and the Causal Structure of the World." In *Scientific Explanation*, 410–505. Minneapolis: University of Minnesota Press.
- . 1992. "The Naturalists Return." *Philosophical Review* 101 (1): 53–114.
- . 2001. *Science, Truth, and Democracy*. New York: Oxford UP.
- Klausen, Søren Harnow. 2009. "Two Notions of Epistemic Normativity." *Theoria* 75: 161–178.
- Knobe, Joshua. 2007. "Experimental Philosophy and Philosophical Significance." *Philosophical Explorations* 10 (2): 119–121.
- , and Shaun Nichols. 2007. "An Experimental Philosophy Manifesto." In *Experimental Philosophy*, edited by Joshua Knobe and Shaun Nichols, 3–14. Oxford: Oxford UP.
- Koertge, Noretta. 1997. "Wrestling with the Social Constructor." In *The Flight from Science and Reason*, edited by Paul Gross, Norman Levitt, and Martin Lewis, 266–273. Baltimore: Johns Hopkins University Press.
- Kornblith, Hilary. 1993. "Epistemic Normativity." *Synthese* 94: 357–376.
- . 2014. *A Naturalistic Epistemology*. Oxford: Oxford UP.

- Krauss, Lawrence. 2012. *A Universe from Nothing: Why There Is Something Rather Than Nothing*. New York: Free Press.
- Kriegel, Uriah. 2013. "The Epistemological Challenge of Revisionary Metaphysics." *Philosophers' Imprint* 13 (12): 1–30.
- Kripke, Saul. 1980. *Naming and Necessity*. Cambridge: Harvard UP.
- Ladyman, James, Don Ross, and Harold Kincaid, eds. 2013. *Scientific Metaphysics*. Oxford: Oxford UP.
- , Don Ross, David Spurrett, and John Collier. 2007. *Every Thing Must Go: Metaphysics Naturalized*. Oxford: Oxford UP.
- Langsam, Harold. 2008. "Rationality, Justification, and the Internalism/Externalism Debate." *Erkenntnis* 68: 79–101.
- Laudan, Larry. 1981. "The Pseudo-Science of Science?" *Philosophy of the Social Sciences* 11: 173–198.
- . 1990a. "Aimless Epistemology?" *Studies in History and Philosophy of Science* 21: 315–322.
- . 1990b. "Demystifying Underdetermination." In *Scientific Theories*, edited by C. Wade Savage, 267–297. Minnesota: University of Minnesota Press.
- . 1990c. "Normative Naturalism." *Philosophy of Science* 57: 44–59.
- Lehrer, Keith, and Stewart Cohen. 1983. "Justification, Truth, and Coherence." *Synthese* 55 (2): 191–207.
- Leibniz, Gottfried. 1709. *New Essays Concerning Human Understanding*.

- Lenski, Richard, and Robert May. 1994. "The Evolution of Virulence in Parasites and Pathogens: Reconciliation Between Two Competing Hypotheses." *Journal of Theoretical Biology* 3 (7): 253–265.
- Leslie, Sarah-Jane. 2013. "Essence and Natural Kinds: When Science Meets Preschooler Intuition." In *Oxford Studies in Epistemology*, edited by Tamar Szabó Gendler and John Hawthorne, 108–165. Oxford: Oxford UP.
- Lewis, David. 1988. "Statements Partly About Observation." *Philosophical Papers* 17 (1): 1–31.
- Lipton, Peter. 2004. *Inference to the Best Explanation*. 2nd edition. London: Routledge.
- Littlejohn, Clayton. forthcoming. "The Right in the Good: A Defense of Teleological Non-Consequentialism in Epistemology." In *Epistemic Consequentialism*, edited by Kristoffer Ahlstrom-Vij and Jeff Dunn. Oxford: Oxford UP.
- . 2013. "The Russellian Retreat." *Proceedings of the Aristotelian Society* 113 (3 pt 3): 293–320.
- . 2015. "Who Cares What You Accurately Believe?" *Philosophical Perspectives* 29 (1): 217–248.
- Livengood, Jonathan, and Edouard Machery. 2007. "The Folk Probably Don't Think What You Think They Think: Experiments on Causation by Absence." *Midwest Studies in Philosophy* 31 (1): 107–127.
- Loewer, Barry. 1996. "Freedom from Physics: Quantum Mechanics and Free Will." *Philosophical Topics* 24: 92–113.
- . 2012. "Two Accounts of Laws and Time." *Philosophical Studies* 160 (1): 115–137.
- Long, Todd. 2010. "Proper Function Justification and Epistemic Rationality." *Southwest Philosophical Review* 26 (1): 189–195.

- Lowe, E.J. 1998. *The Possibility of Metaphysics: Substance, Identity, and Time*. Oxford: Oxford UP.
- Lynch, Michael. 2004. *True to Life: Why Truth Matters*. Cambridge, MA: MIT Press.
- Machery, Edouard, Ron Mallon, Shaun Nichols, and Stephen Stich. 2004. "Semantics, Cross-Cultural Style." *Cognition* 92: B1–B12.
- , and Stephen Stich. 2012. "The Role of Experiment in the Philosophy of Language." In *Routledge Companion to Philosophy of Language*, edited by Gillian Russell and Delia Graff Fara, 495–512. New York: Routledge.
- MacKay, David. 2003. *Information Theory, Inference and Learning Algorithms*. Cambridge: Cambridge UP.
- Mackie, Penelope. 2006. *How Things Might Have Been: Individuals, Kinds, and Essential Properties*. Oxford: Oxford UP.
- Maclaurin, James, and Heather Dyke. 2012. "What Is Analytic Metaphysics For?" *Australasian Journal of Philosophy* 90 (2): 291–306.
- Maddy, Penelope. 2007. *Second Philosophy: A Naturalistic Method*. Oxford: Oxford UP.
- Malcolm, Norman. 1960. "Anselm's Ontological Arguments." *Philosophical Review* 69 (1): 41–62.
- Margenau, Henry. 1967. "Quantum Mechanics, Free Will, and Determinism." *Journal of Philosophy* 64 (21): 714–725.
- Maudlin, Tim. 2007. *The Metaphysics Within Physics*. Oxford: Oxford UP.
- Mayr, Ernst. 1963. *Animal Species and Evolution*. Cambridge, MA: Harvard UP.
- . 1970. *Populations, Species and Evolution*. Cambridge, MA: Harvard UP.

- . 2006. “Typological versus Populational Thinking.” In *Conceptual Issues in Evolutionary Biology*, edited by Elliott Sober, 3rd edition, 325–328. Cambridge, MA: The MIT Press.
- McHugh, Conor. 2011. “What Do We Aim At When We Believe?” *Dialectica* 65 (3): 369–392.
- . 2012. “Epistemic Deontology and Voluntariness.” *Erkenntnis* 77 (1): 65–94.
- McLeod, Mike, and Josh Parsons. 2013. “Maclaurin and Dyke on Analytic Metaphysics.” *Australasian Journal of Philosophy* 91 (1): 173–178.
- Menzies, Peter. 1998. “Possibility and Conceivability: A Response-Dependent Account of Their Connections.” In *European Review of Philosophy, Volume 3: Response-Dependence*, edited by Roberto Casati, 255–277. Stanford: CSLI Publications.
- Mitchell, Sandra. 2002. “Integrative Pluralism.” *Biology and Philosophy* 17: 55–70.
- . 2003. *Biological Complexity and Integrative Pluralism*. Cambridge: Cambridge UP.
- . 2009. *Unsimple Truths: Science, Complexity, and Policy*. Chicago: University of Chicago Press.
- Montmarquet, James. 1993. *Epistemic Virtue and Doxastic Responsibility*. Lanham, MD: Rowman and Littlefield.
- Morganti, Matteo. 2008. “Weak Discernibility, Quantum Mechanics and the Generalist Picture.” *Philosophica* 10 (1/2): 155–183.
- Moser, Paul. 1989. *Knowledge and Evidence*. Cambridge: Cambridge UP.
- Nagel, Thomas. 1986. *The View from Nowhere*. Oxford: Oxford UP.
- Neumann, Manfred. 1978. “A Probabilistic Analysis of the Difficulties of Unifying Quantum Mechanics with the Theory of Relativity.” *Foundations of Physics* 8 (9–10): 721–733.
- Ney, Alyssa. 2012. “The Status of Our Ordinary Three Dimensions in a Quantum Universe.” *Noûs* 46 (3): 525–560.



- , and David Albert, eds. 2013. *The Wave Function: Essays in the Metaphysics of Quantum Mechanics*. Oxford: Oxford UP.
- Nichols, Shaun, and Joshua Knobe. 2007. “Moral Responsibility and Determinism: The Cognitive Science of Folk Intuitions.” *Noûs* 41 (4): 663–685.
- , Shikhar Kumar, Theresa Lopez, Alisabeth Ayars, and Hoi-Yee Chan. 2016. “Rational Learners and Moral Rules.” *Mind and Language* 31 (5): 530–554.
- , Stephen Stich, and Jonathan Weinberg. 2003. “Metaskepticism: Meditations in Ethnoepistemology.” In *The Sceptics*, edited by Steven Luper, 227–247. Burlington: Ashgate Publishing.
- Nolan, Daniel. 2017. “Naturalised Modal Epistemology.” In *Modal Epistemology After Rationalism*, edited by Bob Fischer and Felipe Leon, 7–27. Dordrecht: Springer.
- Norton, Joshua. 2015. “Weak Discernibility and Relations Between Quanta.” *Philosophy of Science* 82 (5): 1188–1199.
- Nozick, Robert. 1981. *Philosophical Explanations*. Cambridge, MA: Harvard UP.
- . 1993. *The Nature of Rationality*. Princeton, NJ: Princeton UP.
- Okasha, Samir. 2002. “Darwinian Metaphysics: Species and the Question of Essentialism.” *Synthese* 131 (2): 191–213.
- Owens, David. 2000. *Reason Without Freedom: The Problem of Epistemic Normativity*. New York: Routledge.
- Papineau, David. 2014. “The Poverty of Conceptual Analysis.” In *Philosophical Methodology: The Armchair or the Laboratory*, edited by Matthew Haug, 166–194. Abingdon: Routledge.

- Park, Leeyoung, and Ju Kim. 2015. "A Novel Approach for Identifying Causal Models of Complex Diseases from Family Data." *Genetics* 199 (4): 1007–1016.
- Peacocke, Chris. 1999. *Being Known*. Oxford: Oxford UP.
- Pennock, Robert. 2000. *Tower of Babel: The Evidence Against the New Creationism*. Cambridge, MA: MIT Press.
- Perfors, Amy, Joshua Tenenbaum, Thomas Griffiths, and Fei Xu. 2011. "A Tutorial Introduction to Bayesian Models of Cognitive Development." *Cognition* 120: 302–321.
- Petersen, Steve. 2013. "Utilitarian Epistemology." *Synthese* 190 (6): 1173–1184.
- Petrinovich, Lewis, and Patricia O'Neill. 1996. "Influence of Wording and Framing Effects on Moral Intuitions." *Ethology and Sociobiology* 17 (3): 145–171.
- Plantinga, Alvin. 1974. *The Nature of Necessity*. Oxford: Oxford UP.
- Pradeu, Thomas, and Alexandre Guay. 2015. *Individuals Across the Sciences*. Oxford: Oxford UP.
- Priest, Graham. 1998. "What Is So Bad About Contradictions?" *Journal of Philosophy* 95 (8): 410–426.
- . 2014. *One: Being an Investigation Into the Unity of Reality and of Its Parts, Including the Singular Object Which Is Nothingness*. Oxford: Oxford UP.
- Prinz, Jesse. 2008. "Empirical Philosophy and Experimental Philosophy." In *Experimental Philosophy*, edited by Joshua Knobe and Shaun Nichols, 189–208. Oxford: Oxford UP.
- Prior, A.N. 1960. "The Autonomy of Ethics." *Australasian Journal of Philosophy* 38 (3): 199–206.
- Psillos, Stathis. 1999. "Empiricism and Theoretical Discourse." In *Scientific Realism: How Science Tracks Truth*, 3–15. New York: Routledge.

- Putnam, Hilary. 1967. "Time and Physical Geometry." *Journal of Philosophy* 64: 240–247.
- . 1973. "Meaning and Reference." *Journal of Philosophy* 70: 699–711.
- . 1982. "Why Reason Can't Be Naturalized." *Synthese* 52 (1, Realism, Part II): 3–23.
- Quine, Willard van Orman. 1951. "Two Dogmas of Empiricism." *Philosophical Review*.  
Cambridge, Mass: Harvard UP, 20–46.
- . 1969. "Epistemology Naturalized." In *Ontological Relativity and Other Essays*. New York: Columbia UP.
- . 1981. *Theories and Things*. Cambridge, MA: Harvard UP.
- Resnik, David. 2000. "A Pragmatic Approach to the Demarcation Problem." *Studies in History and Philosophy of Science Part A* 31 (2): 249–267.
- Rey, Georges. 1983. "Concepts and Stereotypes." *Cognition* 15 (1–3): 237–262.
- . 1985. "Concepts and Conceptions: A Reply to Smith, Medin and Rips." *Cognition* 19 (3): 297–303.
- Roca-Royes, Sonia. 2007. "Mind-Independence and Modal Empiricism." In *4th Latin Meeting in Analytic Philosophy, CEUR-WS Proceedings*, edited by Carlo Penco, Massimiliano Vignolo, Valeria Ottonelli, and Cristina Amoretti, 117–135. Genoa: University of Genoa.
- . 2011. "Conceivability and De Re Modal Knowledge." *Noûs* 45 (1): 22–49.
- Rose, David, and Jonathan Schaffer. 2015. "Folk Mereology Is Teleological." *Noûs* 49 (3).
- Rosen, Gideon. 2010. "Metaphysical Dependence: Grounding and Reduction." In *Modality: Metaphysics, Logic, and Epistemology*, edited by Robert Hale and Aviv Hoffman, 109–136. New York: Oxford UP.

- Rosenberg, Alexander. 1994. *Instrumental Biology, Or, The Disunity of Science*. Chicago: University of Chicago Press.
- Ruetsche, Laura. 2011. *Interpreting Quantum Theories: The Art of the Possible*. Oxford: Oxford UP.
- Ruphy, Stéphanie. 2017. *Scientific Pluralism Reconsidered*. Pittsburgh: Pittsburgh UP.
- Ryan, Sharon. 2003. "Doxastic Compatibilism and the Ethics of Belief." *Philosophical Studies* 114: 47–79.
- Salmon, Wesley. 1984. *Scientific Explanation and the Causal Structure of the World*. Princeton: Princeton UP.
- Schaffer, Jonathan. 2003. "Is There a Fundamental Level." *Noûs* 37 (3): 498–517.
- . 2009. "On What Grounds What." In *Metametaphysics: New Essays on the Foundations of Ontology*, edited by David Manley, David Chalmers, and Ryan Wasserman, 347–383. Oxford: Oxford UP.
- . 2010. "Monism: The Priority of the Whole." *Philosophical Review* 119 (1): 31–76.
- . 2012. "Grounding, Transitivity, and Contrastivity." In *Metaphysical Grounding: Understanding the Structure of Reality*, edited by Fabrice Correia and Benjamin Schneider, 122–138. Cambridge: Cambridge UP.
- . 2015. "Grounding in the Image of Causation." *Philosophical Studies Online*. [www.jonathanschaffer.org/groundcause.pdf](http://www.jonathanschaffer.org/groundcause.pdf).
- Schneider, Benjamin, Susan White, and Michelle Paul. 1998. "Linking Service Climate and Customer Perceptions of Service Quality: Test of a Causal Model." *Journal of Applied Psychology* 83 (2): 150–163.

- Schwitzgebel, Eric, and Fiery Cushman. 2012. "Expertise in Moral Reasoning? Order Effects on Moral Judgment in Professional Philosophers and Non-Philosophers." *Mind and Language* 27 (2): 135–153.
- Shaffer, Michael. 2013. "Doxastic Voluntarism, Epistemic Deontology and Belief-Contravening Commitments." *American Philosophical Quarterly* 50 (1): 73–82.
- Shoemaker, Sydney. 1980. "Causality and Properties." In *Time and Cause*, edited by Peter Van Inwagen, 109–135. Dordrecht: Reidel Publishing Co.
- . 1998. "Causal and Metaphysical Necessity." *Philosophical Quarterly* 79 (1): 59–77.
- Shtulman, Andrew, and Kelsey Harrington. 2016. "Tensions Between Science and Across the Lifespan." *Topics in Cognitive Science* 8: 118–137.
- Singer, Daniel. 2015. "Mind the Is-Ought Gap." *Journal of Philosophy* 112 (4): 193–210.
- Sinnott-Armstrong, Walter. 2008. "Framing Moral Intuitions." In *Moral Psychology, Vol. 2: The Cognitive Science of Morality*, edited by Walter Sinnott-Armstrong, 47–76. Cambridge, MA: MIT Press.
- Sklar, Lawrence. 1974. *Space, Time, and Spacetime*. Berkeley: University of California Press.
- Smithies, Declan. 2012. "The Normative Role of Knowledge." *Noûs* 46 (2): 265–288.
- Smolin, Lee. 2006. *The Trouble with Physics: The Rise of String Theory, The Fall of a Science, and What Comes Next*. New York: Houghton Mifflin.
- Sober, Elliott. 1980. "Evolution, Population Thinking and Essentialism." *Philosophy of Science* 47: 350–383.
- . 1993. *Philosophy of Biology*. Boulder: Westview Press.
- . 2015. *Ockham's Razors: A User's Manual*. Cambridge: Cambridge UP.
- Sosa, Ernest. 1980. "The Foundations of Foundationalism." *Noûs* 14 (4): 547–564.

- . 2007. *A Virtue Epistemology: Apt Belief and Reflective Knowledge, Volume I*. Oxford: Oxford UP.
- Stapleford, Scott. 2015. “Why There May Be Epistemic Duties.” *Dialogue* 54 (1): 63–89.
- Steglich-Petersen, Asbjørn. 2013. “Truth as the Aim of Epistemic Justification.” In *The Aim of Belief*, edited by Timothy Chan, 204–226. Oxford: Oxford UP.
- Stein, Howard. 1968. “On Einstein-Minkowski Space-Time.” *Journal of Philosophy* 65: 5–23.
- . 1970. “A Note on Time and Relativity Theory.” *Journal of Philosophy* 67: 289–294.
- Stern, Robert. 2009. *Hegelian Metaphysics*. Oxford: Oxford UP.
- Steup, Matthias. 2000. “Doxastic Voluntarism and Epistemic Deontology.” *Acta Analytica* 15 (1): 25–56.
- . 2008. “Doxastic Freedom.” *Synthese* 161: 375–392.
- Stich, Stephen. 1993. “Naturalizing Epistemology: Quine, Simon and the Prospects for Pragmatism.” In *Philosophy and Cognitive Science, Royal Institute of Philosophy, Supplement No. 34*, edited by Christopher Hookway and Donald Peterson, 1–17. Cambridge: Cambridge UP.
- Strohming, Margot. 2015. “Perceptual Knowledge of Nonactual Possibilities.” *Philosophical Perspectives* 29: 363–375.
- Sutton, Jonathan. 2007. *Without Justification*. Cambridge, MA: MIT Press.
- Swain, Stacey, Joshua Alexander, and Jonathan Weinberg. 2008. “The Instability of Philosophical Intuitions: Running Hot and Cold on Truetemp.” *Philosophy and Phenomenological Research* 76 (1): 138–155.
- Swinburne, Richard. 2001. *Epistemic Justification*. Oxford: Oxford UP.

- Swoyer, Chris. 1982. "The Nature of Natural Laws." *Australasian Journal of Philosophy* 60: 203–223.
- Sylvan, Kurt. forthcoming. "Veritism Unswamped." *Mind*.
- Sytsma, Justin, and Jonathan Livengood. 2015. *The Theory and Practice of Experimental Philosophy*. Peterborough: Broadview Press.
- , and Edouard Machery. 2009. "How to Study Folk Intuitions About Phenomenal Consciousness." *Philosophical Psychology* 22 (1): 21–35.
- Tenenbaum, Joshua, and Thomas Griffiths. 2001. "Generalization, Similarity, and Bayesian Inference." *Behavioral and Brain Sciences* 24: 629–640.
- Thomas, Alan. 2008. "The Genealogy of Epistemic Virtue Concepts." *Philosophical Papers* 37 (3): 345–369.
- Thomson, Judith Jarvis. 1998. "The Statue and the Clay." *Noûs* 32 (2): 149–173.
- Vahid, Hamid. 2003. "Truth and the Aim of Epistemic Justification." *Teorema* 20 (3): 83–91.
- Vaidya, Anand. 2015. "The Epistemology of Modality." In *The Stanford Encyclopedia of Philosophy*, edited by Edward Zalta.  
<http://plato.stanford.edu/archives/sum2015/entries/modality-epistemology/>.
- Van Fraassen, Bas. 1980. *The Scientific Image*. Oxford: Oxford UP.
- van Inwagen, Peter. 1981. "The Doctrine of Arbitrary Undetached Parts." *Pacific Philosophical Quarterly* 62 (April): 123–137.
- . 1998. "Modal Epistemology." *Philosophical Studies* 92: 67–84.
- Velleman, David. 2000. *The Possibility of Practical Reason*. Oxford: Oxford UP.
- Vienna Circle, The. 1929. "The Scientific Conception of the World."  
[evidencebasedcryonics.org/pdfs/viennacircle.pdf](http://evidencebasedcryonics.org/pdfs/viennacircle.pdf).

- Waters, Ken. 2014. "Shifting Attention From Theory to Practice in Philosophy of Biology." In *New Directions in the Philosophy of Science*, edited by Maria Carla Galavotti, Dennis Dieks, Wenceslao J. Gonzalez, Stephan Hartmann, Thomas Uebel, and Marcel Weber, 121–139. Berlin: Springer International Publishing.
- . 2017. "No General Structure." In *Metaphysics and the Philosophy of Science: New Essays*, edited by Matthew Slater and Zanja Yudell, 81–108. Oxford: Oxford UP.
- Weatherson, Brian. 2008. "Deontology and Descartes' Demon." *The Journal of Philosophy* 150: 540–569.
- Weinberg, Jonathan, Chad Gonnerman, Cameron Buckner, and Joshua Alexander. 2010. "Are Philosophers Expert Intuiters?" *Philosophical Psychology* 23 (2): 331–355.
- , Shaun Nichols, and Stephen Stich. 2001. "Normativity and Epistemic Intuitions." *Philosophical Topics* 29 (1–2): 429–460.
- Weiner, Matthew. 2005. "Why Does Justification Matter?" *Pacific Philosophical Quarterly* 86: 422–444.
- . 2014. "The Spectra of Epistemic Norms." In *Epistemic Norms: New Essays on Action, Belief, and Assertion*, edited by Clayton Littlejohn and John Turri, 201–218. Oxford: Oxford UP.
- Weingard, Robert. 1972. "Relativity and the Reality of Past and Future Events." *British Journal for the Philosophy of Science* 23: 119–121.
- Wheatley, Thalia, and Jonathan Haidt. 2005. "Hypnotic Disgust Makes Moral Judgments More Severe." *Psychological Science* 16 (10): 780–784.
- Whiting, Daniel. 2012. "Does Belief Aim at the Truth?" *Pacific Philosophical Quarterly* 93 (2): 279–300.



- . 2013a. “Nothing but the Truth: On the Norms and Aims of Belief.” In *The Aim of Belief*, edited by Timothy Chan, 184–203. Oxford: Oxford UP.
- . 2013b. “The Good and the True (or the Bad and the False).” *Philosophy* 8 (2): 219–242.
- . 2013c. “Truth: The Aim and Norm of Belief.” *Teorema* 32 (3): 121–136.
- Wiegmann, Alex, Yasmina Okan, and Jonas Nagel. 2012. “Order Effects in Moral Judgment.” *Philosophical Psychology* 25 (6): 813–836.
- Williamson, Timothy. 2000. *Knowledge and Its Limits*. Oxford: Oxford UP.
- . 2007. “Knowledge of Metaphysical Modality.” In *The Philosophy of Philosophy*, 134–179. Malden, MA: Blackwell.
- . 2014. “No Work for a Theory of Grounding.” *Inquiry* 57 (5–6): 535–579.
- Wilson, Mark. 2006. *Wandering Significance: An Essay on Conceptual Behaviour*. Oxford: Oxford UP.
- Wimsatt, William. 1994. “The Ontology of Complex Systems: Levels of Organization, Perspectives, and Causal Thickets.” *Canadian Journal of Philosophy* 20 (sup 1): 207–274.
- Woit, Peter. 2006. *Not Even Wrong: The Failure of String Theory and the Search for Unity in Physical Law*. New York: Basic Books.
- Wolchover, Natalie. 2015. “Physicists and Philosophers Hold Peace Talks.” *The Atlantic*. [www.theatlantic.com/science/archive/2015/12/physics-philosophy-string-theory/421569/?utm\\_source=SFFB](http://www.theatlantic.com/science/archive/2015/12/physics-philosophy-string-theory/421569/?utm_source=SFFB).
- Worley, Sara. 2003. “Conceivability, Possibility and Physicalism.” *Analysis* 63: 15–23.

- Wright, Sarah. 2014. "The Dual-Aspect Norms of Belief and Assertion: A Virtue Approach to Epistemic Norms." In *Epistemic Norms: New Essays on Action, Belief, and Assertion*, edited by Clayton Littlejohn and John Turri, 239–258. Oxford: Oxford UP.
- Yablo, Stephen. 1993. "Is Conceivability a Guide to Possibility?" *Philosophy and Phenomenological Research* 53 (1): 1–42.
- Zagzebski, Linda. 1996. *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*. Cambridge: Cambridge UP.
- . 2003. "Intellectual Motivation and the Good of Truth." In *Intellectual Virtue: Perspectives from Ethics and Epistemology*, edited by Linda Zagzebski and Michael DePaul, 135–154. New York: Oxford UP.
- , and Abrol Fairweather, eds. 2000. *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility*. Oxford: Oxford UP.