

Empiricism, Natural Regularity, and Necessity

by

Tyler William Hildebrand

B.A., Pacific Lutheran University, 2006

M.A., University of Colorado at Boulder, 2009

A thesis submitted to the
Faculty of the Graduate School of the
University of Colorado in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
Department of Philosophy

2011

This thesis entitled:
Empiricism, Natural Regularity, and Necessity
written by Tyler William Hildebrand
has been approved for the Department of Philosophy

Professor Michael Tooley

Professor Graeme Forbes

Professor Graham Oddie

Professor Robert Hanna

Date _____

The final copy of this thesis has been examined by the signatories, and we find that both the content and the form meet acceptable presentation standards of scholarly work in the above mentioned discipline.

Hildebrand, Tyler William (Ph.D., Philosophy)

Empiricism, Natural Regularity, and Necessity

Thesis directed by Professor Michael Tooley

This dissertation has two parts. In the first, I set out and defend a new empirical method of practicing metaphysics. The method avoids appeals to intuitions, ordinary beliefs, and the like. It does not accept basic principles of simplicity, unity, and the like. Instead, it proceeds from logic, analytic principles, and immediate experience alone. In the second part of my dissertation, I apply this method to the philosophy of laws of nature. I argue that there are excellent empirical reasons to accept governing laws instead of laws that reduce to other features of the world, such as natural regularities or facts about bare dispositions. The central idea is that observed natural regularities constitute strong evidence in favor of governing laws and against all competing theories. Further, I argue that the only intelligible account of governing laws is one according to which the connection between law and regularity is an irreducible necessary connection. Thus, the second part of my dissertation constitutes a new argument for metaphysically interesting a posteriori necessities.

Contents

Chapter

1	Introduction	1
	1.1 Introduction	1
	1.2 Natural Regularities	2
	1.3 Theories of Laws	3
	1.4 Chapter Outlines	7
2	Methodology	10
	2.1 Introduction	10
	2.2 Epistemological Assumptions	11
	2.3 An Argument Against Empirical Metaphysics	13
	2.4 Empirical Resources	16
	2.4.1 Empirical Foundationalism	16
	2.4.2 Laws as Theoretical Entities	17
	2.4.3 The Method of Confirmation	26
	2.4.4 Assigning Probabilities	29
	2.5 The Method	44
	2.6 Conclusion	48
	2.7 Appendix: Miscellaneous Methodological Criteria	49
	2.7.1 Ontological Economy and Occam's Razor	49
	2.7.2 The Status of Modal Intuitions	52

3	Non-Humeanism	54
3.1	Introduction	54
3.2	A Theoretical Analysis of Synthetic Necessity	55
3.2.1	The Analysis of Metaphysical Possibility	55
3.2.2	More on the Concept of Actual Restricted Interpretations	61
3.2.3	Synthetic A Priori Reasoning and Restricted Interpretations	63
3.2.4	An A Priori Argument Against Humeanism?	63
3.2.5	Implications and Further Discussion	64
3.3	Bare Dispositionalism: A Type of Non-Humeanism	66
3.3.1	The Dispositional / Categorical Distinction and Realism vs. Reductionism	67
3.3.2	Three Arguments Against Bare Dispositionalism	69
3.3.3	Summing Up	76
3.4	The A Priori Probabilities of Descriptive Humeanism and Governing Non-Humeanism	77
3.5	Next Steps	84
4	Governing Humeanism	85
4.1	Introduction	85
4.2	Tooley's Version of Governing Humeanism	86
4.2.1	Preliminaries	86
4.2.2	Tooley's Speculative Account	89
4.2.3	The Objection from Probabilistic Laws	92
4.2.4	The Objection from Temporally Extended Laws	93
4.2.5	An Objection Concerning the Definition of Conjunctive Universals	97
4.3	Contemporary Arguments for the Inconsistency of Governing Humeanism	99
4.3.1	van Fraassen's Argument	99
4.3.2	Bird's Argument	102
4.4	A Modified Argument for the Inconsistency of Governing Humeanism	103

4.5	Conclusion	107
5	Descriptive Humeanism	109
5.1	Introduction	109
5.2	Preliminaries	110
5.2.1	Regularity	110
5.2.2	Review of Assumptions	111
5.3	Previous Explanatory Arguments against Descriptive Humeanism	112
5.3.1	David Armstrong's Argument	112
5.3.2	Evan Fales's Argument	115
5.4	The Argument from General Regularity	118
5.4.1	In Contrast with the Previous Arguments	118
5.4.2	A Statement of the General Argument	120
5.4.3	Equiprobability of State Descriptions	122
5.4.4	Equiprobability of Structure Descriptions	133
5.5	Objections	136
5.6	A Big Picture Objection	138
5.6.1	Probabilistic Laws	139
5.6.2	Deterministic Laws	142
5.7	Conclusion	145
6	Descriptive Non-Humeanism	146
6.1	Introduction	146
6.2	The A Priori Probabilities	148
6.3	How Descriptive Non-Humeanism Explains Regularities	149
6.4	The Nature of Observed Regularities	151
6.4.1	The Argument From Experience	152
6.4.2	The Regress Argument Against Dispositional Monism	156

6.4.3	Does Dualism Imply Quidditism? Is That Bad?	160
6.5	The Explanatory Weakness of Descriptive Non-Humeanism	165
6.6	Conclusion: The Explanatory Power of Governing Non-Humeanism	169
7	Conclusion	172
7.1	Final Remarks on Laws of Nature	172
7.2	Final Remarks on Method	175
7.2.1	Relation to Quine's Empirical Method	175
7.2.2	Contemporary Debates in Metaontology	178
	Bibliography	181

Tables

Table

3.1	Humean State Descriptions for $L: \Omega = \emptyset$	78
3.2	Non-Humean State Descriptions for $L: \Omega = \{F \rightarrow G\}$	79
3.3	Non-Humean State Descriptions for $L: \Omega = \{G \rightarrow F\}$	79
3.4	Non-Humean State Descriptions for $L: \Omega = \{F \rightarrow G, G \rightarrow F\}$	80
5.1	Humean State Descriptions for $L: \Omega = \emptyset$	126
5.2	Non-Humean State Descriptions for $L: \Omega = \{F \rightarrow G\}$	127
5.3	Non-Humean State Descriptions for $L: \Omega = \{G \rightarrow F\}$	127
5.4	Non-Humean State Descriptions for $L: \Omega = \{F \rightarrow G, G \rightarrow F\}$	128

Chapter 1

Introduction

1.1 Introduction

Natural necessities, if there are any, are special. They are unlike normal empirical facts in that they *must* (in at least some sense of ‘must’) be true. Compare ‘the stone must fall if unsus-pended’ with ‘the stone is heavy’. They are unlike the propositions of mathematics and logic in that they are conceivably false. We can imagine the stone to float in midair without contradiction; we cannot imagine $2 + 2 = 4$ to be false, no matter our effort. In the case of mathematics and logic, we (or at least most philosophers) find the necessity involved to be clear and accessible; we don’t find serious epistemic problems in holding such propositions to be necessarily true. The same cannot be said for natural necessity.¹ In this case, we can’t “see” that the propositions are true just by thinking about them or by examining their meaning or logical form. Something else is required. Many philosophers have thought that this something else can’t possibly be *empirical*. I shall argue that they are mistaken.

This project has two basic parts. In the first part, I set forth and defend an empirical method of practicing metaphysics. The method does not appeal to intuitions, ordinary beliefs, and the like. It does not accept basic principles of simplicity, unity, and the like. Instead, I argue that the practice of metaphysics can proceed from logic, analytic principles, and immediate experience alone. I believe that the defense of this method is the most important contribution of this project,

¹ For the moment, I am not assuming that natural necessities are distinct from metaphysical necessities, though perhaps they are. What is important for the moment is that natural necessities are *synthetic* rather than *analytic*.

but I won't discuss it in any detail in this chapter. Instead, the purpose of this chapter is to explain my very general approach to thinking about the philosophy of laws of nature. The reason for this is that metaontological problems are most easily explained in the context of a particular issue in ontology. The chapter provides that context.

In the second part of this project, I apply the method to argue that observed natural regularities constitute decisive evidence for governing laws of nature. Further, I argue that governing laws must be understood as involving irreducible synthetic necessities. Hence, my argument for governing laws of nature constitutes a new argument for metaphysically interesting a posteriori necessities.

1.2 Natural Regularities

As one with empiricist inclinations, I believe that our access to laws of nature begins with our observations of the natural world. Our world has a very interesting feature: it is full of *natural regularities*, which is to say (roughly) that our experiences of the world can be systematized using general principles, that types of experiences occur in certain repeated patterns.² For example, every rock we have ever observed here on Earth falls to the ground when it is unsuspending. Stones dropped into ponds always create a wave pattern of concentric circles. We have observed regularities as far as observation takes us, at both macro and micro levels. Of course, some regularities are not immediately obvious. It took us quite some time to recognize the correlation between genes and macroscopic physical traits of living organisms. And sometimes we are mistaken about particular regularities. For a long time, everyone believed that the universe was Newtonian. The Eddington experiments—which revealed a particular *irregularity* from the perspective of the Newtonian theory—gave us reason to doubt this theory. Newton's theory was widely rejected, but the view that the world is regular was not. Instead, a new theory, Einstein's theory of relativity, took its place, and allowed us to better systematize our observations of particular facts into a general system. Though Newton's theory was rejected due to the fact that some of the regularities it pos-

² I shall offer a more careful definition of natural regularity in Chapter 5.

tulated were false, we did not abandon our belief that the world is regular, that our world can be systematized using a set of general principles.

My basic approach to solving philosophical problems concerning laws of nature is to focus on natural regularities. After all, the feature of the world that makes the process of systematization possible, regularity, is somewhat incredible. If the repeated patterns or generalizations weren't finite, or if they were too large, we could not systematize them. If the world weren't regular at all, we could not systematize our observations into any sort of useful system. Thus I take natural regularities to be the natural starting point for any discussion on laws of nature. They are our basic evidence, our epistemic contact with laws (or whatever it is that explains or is responsible for regularities). In fact, I believe that regularities constitute strong evidence in favor of some theories of laws and against others. This evidence has been neglected by most philosophers working in the area of laws of nature.

1.3 Theories of Laws

Recently, most of the debate concerning laws of nature has focused on five rather narrow theories. The first is the Dretske/Tooley/Armstrong theory according to which laws are atomic states of affairs, consisting of second-order external relations between first-order universals, that entail or probabilify regularities (see (Dretske 1977), (Tooley 1977), and (Armstrong 1983)). The second is John Carroll's (1994) view that laws, though they are states of affairs distinct from regularities that entail or probabilify regularities, cannot be given any further analysis. The third is that laws reduce to or supervene on regularities in a Humean base (one that, at a minimum, does not involve any synthetic necessary connections). This is Hume's view, and has recently been defended by (Lewis 1994), (Loewer 1996), (Earman and Roberts 2005a), (Earman and Roberts 2005b), (Schaffer 2008), and (Beebe 2000). The fourth is that laws reduce to regularities of a very special sort: namely, regularities over a Non-Humean base in which (at least some) natural properties are taken to be bare dispositions, properties with a special sort of intrinsic modal character. (Bird 2007) and (Ellis 2001) have recently defended versions of this view. The fifth, defended by

(van Fraassen 1989), is that there simply are no laws. One way to get a clearer understanding of laws of nature is to carefully consider each of these competitors and determine which is best. But this method is problematic in certain respects. For one, it tends to encourage arguments that are very narrow in scope, applying only to the narrow theories above. Though not a philosophical problem, this makes the dialectic of the philosophy of laws of nature somewhat difficult to follow. For another, the theories explained above are not exhaustive. There are other possible theories of laws, some of which deserve careful attention. As we shall see in the next chapter, this situation is problematic because it makes genuine confirmation of such theories in light of relevant evidence difficult, if not impossible.

Fortunately, these problems can be remedied rather easily. One simply has to offer general definitions of theories that are mutually exclusive and jointly exhaustive. I prefer to do so as follows. Every theory of laws of nature can be placed into one of four categories generated by two distinctions. These distinctions have been chosen carefully, though the reasons for drawing them in this way will not become fully clear until later chapters. The first crucial distinction is between Humeanism and Non-Humeanism.

Humeanism: a proposition is necessarily true (false) if and only if it is true (false) in virtue of its meaning or logical form (that is, if and only if it is *analytically* true (false)).

Non-Humeanism: Humeanism is false; some synthetic propositions are necessarily true.

Humeanism entails that there are no *synthetic necessities*. The Humean has no trouble explaining why it is a necessary truth that $Fa \& Ga$ entails Fa . This is a simple case of logical entailment. However, suppose now that $a \neq b$, $F \neq G$, and F and G refer to logically distinct properties. The Humean will find the suggestion that Fa entails Gb unintelligible.³ This is the sort of necessary connection that Humeanism denies but Non-Humeanism (may) allow; it is a basic necessity, not analyzable in terms of meaning or logical form. Thus, Humeanism accepts, whereas Non-Humeanism rejects,

³ I am, for the moment, ignoring Tooley's (1987) speculative explanation of the entailment that holds between law and regularity. However, in Chapter 4 I shall argue that Humeanism is incompatible with the thesis that laws govern.

the following dichotomy: all and only necessary propositions are analytic; all and only contingent propositions are synthetic. Ultimately, I shall be defending a version of Non-Humeanism.

(I used to think that Non-Humeanism could be defined in terms of *de re* necessities. I now think that this way of drawing the distinction is problematic. While the *de re / de dicto* distinction is both interesting (see (Fine 1978)) and relevant to our purposes, it is too narrow. There are *de dicto* propositions that are unacceptable from the Humean perspective, such as the proposition that *necessarily, nothing travels faster than the speed of light*. The Humean will be skeptical of this proposition despite the fact that it is *de dicto*. It is synthetic. No necessary connection can be discovered between the concepts of velocity and light; there is no relevant relation between ideas. The necessity must be taken as a brute, bare, irreducible necessity. Hume did not like such necessities, and thus it would be misleading to characterize any view that accepted them as ‘Humean’.)

The second crucial distinction is between *Governing* and *Descriptive* (or *Non-Governing*) theories of laws. A very rough way to draw this distinction is as follows: according to a descriptive theory of laws, laws (if there are any) merely describe the world, so the laws depend on the regularities they describe; according to a governing theory of laws, however, laws actively “shape” the world, so the regularities depend on the laws. More precisely:

Governing Laws: (a) Laws are states of affairs distinct from regularities that entail (or probabilify) regularities, and (b) there is at least one law.

Descriptive Laws: There are no governing laws; if there are any laws at all, they reduce to or supervene on other features of the world (such as facts about regularities or bare dispositions).

Of course, these are not the only distinctions relevant to laws of nature, but these are the distinctions relevant to the matter of whether there are any synthetic necessities. Further distinctions may be easily introduced later within the context of the four available types of theories generated by the above two distinctions, but these four theories are the primary focus of this project. I shall now offer a quick explanation of the four possible theories of laws of nature and situate them within recent literature.

Descriptive Humeanism accepts Humeanism and Descriptive Laws. The view has two prominent versions: those theories that hold that there are laws (*reductionism*), and those theories that hold that there are not (*anti-realism*). David Hume, David Lewis (1973) and (1994), van Fraassen (1989), Loewer (1996), and Earman and Roberts (2005a) and (2005b) are all Descriptive Humeans. They all accept a Humean ontology and reject the thesis that laws govern.

Governing Humeanism accepts Humeanism and Governing Laws. This is the view of Armstrong (1983), Tooley (1987), and Dretske (1977).⁴ Ultimately, I shall argue that Governing Humeanism is internally inconsistent on the grounds that the necessary connection between law and regularity must be taken as primitive, and is thus incompatible with Humeanism.

Descriptive Non-Humeanism accepts Non-Humeanism and Descriptive Laws. Like its Humean counterpart, it comes in both *reductionist* ((Fales 1990), (Ellis 2001) (perhaps), and (Bird 2007)) and *anti-realist* (Mumford 2004) varieties. Versions of such theories are often called *dispositional essentialism* or *necessitarianism* since they generally hold that the modal properties that ground laws—that is, dispositional properties—are dispositions essential to the objects that instantiate them, and they have the implication that laws of nature are metaphysically necessary.

Finally, there is the view that I prefer, *Governing Non-Humeanism*, which accepts Non-Humeanism and Governing Laws. The entailment between law and regularity is analyzed as a synthetic necessary connection. There are different ways that one could accept such a view. One could hold that the analysis of laws make reference to some particular. For example, Isaac Newton and George Berkeley both thought that God was in some sense responsible for the laws of nature. The idea here is that a law is a sort of preference in the mind of God that the world be thus and so; that preference is a state of affairs distinct from regularities, and (for some preferences) it entails regularities. Alternatively, it may be that the world as a whole possesses a grand disposition that is responsible for the regularities of the world (this is one way of interpreting Ellis's (2001) own view). Alternatively, one could hold that there is an analysis of the concept of laws that makes no

⁴ It is less clear that Dretske accepts Humeanism, though his (1977) claim that nomic necessitation is “an extensional relation” suggests this reading. Since he is not as explicit in his endorsement of Humeanism, I shall usually just attribute the Governing Humean view to Tooley and Armstrong.

reference to particulars; this would be the case if one followed Dretske, Tooley, and Armstrong in thinking that laws are irreducible relations between universals (but did not follow them in thinking that the connection involves no irreducible necessity). Finally, one could hold that there is no analysis of the concept of laws beyond the idea that laws are states of affairs distinct from regularities that entail (or probabilify) regularities (this is more or less Carroll's (1994) view). For the most part I shall ignore these narrower distinctions.

As mentioned earlier, the distinctions between the four general theories of laws are relevant to the problem of whether there are any synthetic necessities. The distinction between Humeanism and Non-Humeanism merely captures the two sides of this problem. The distinction between Governing and Descriptive theories is relevant because some philosophers (Tooley and Armstrong) have claimed that their Humean theories (remember, they are Humeans in that they reject synthetic necessity, but they are not Humeans about laws in the traditional sense of the word) have explanatory power with respect to the regularities of the natural world; that is, they claim that the probability of a world full of regularities given Governing Humeanism is just as high as the probability of a world full of regularities given Governing Non-Humeanism. If their claim is correct, then the matter of using natural regularities as evidence in favor of Non-Humeanism becomes significantly more complicated and controversial. Thus the matter of whether there are any synthetic necessities cannot (in the context of my argument) be detached from the metaphysics of laws of nature. Ultimately, my thesis is that Governing Non-Humeanism is preferable to all of the alternatives.

1.4 Chapter Outlines

In Chapter 2, I explain my basic approach to solving metaphysical problems. I assume a version of empiricism according to which all descriptive concepts are formed on the basis of the immediately observable, and according to which the justification of synthetic principles is ultimately grounded in the immediately observable. I argue that nomological concepts require theoretical analyses. I then explain the following three technical apparatuses: the Ramsey/Lewis method of

defining theoretical terms, a probabilistic method of empirical confirmation, and a version of Carnap's logical interpretation of probability. Finally, I explain how a properly constrained synthesis of these methods allows for the practice of genuine metaphysics. In short, I argue that the practice of metaphysics can proceed from immediate observation and analytic principles alone.

In Chapter 3, I offer a preliminary defense of Non-Humeanism by arguing that there is nothing problematic with the idea of irreducible synthetic necessity. I do so by offering an analysis of the sort of synthetic necessity characteristic of Non-Humeanism that is compatible with the version of empiricism I set out in Chapter 2 (it is a very "thin" analysis, capturing both a priori and a posteriori necessity). I then argue that, given the intelligibility of Non-Humeanism, the intrinsic character of Governing Non-Humeanism makes it a priori just as likely as Descriptive Humeanism. In short, I argue that there is no reason (even for the empiricist) to prejudice Governing Non-Humeanism on the grounds that it incorporates synthetic necessity.

In Chapter 4, I consider Tooley's Humean-consistent account of the necessary connection between governing law and regularity and argue that it fails. I then argue that Governing Humeanism is self-contradictory. The argument is based on the idea that it is impossible to offer a Humean analysis of the entailment that holds between governing laws and regularities. The only possible necessary connection between the two is synthetic. Thus, all versions of Humeanism are committed to Descriptive Laws. The result is that Governing Non-Humeanism is superior to Governing Humeanism.

In Chapter 5, I provide the basic empirical argument against Descriptive Humeanism. This argument is based on the fact that Descriptive Humeanism, unlike Governing Non-Humeanism, lacks explanatory power with respect to the regularities of the natural world. Since the conclusion of Chapter 3 is that Governing Non-Humeanism is a priori just as likely as Descriptive Humeanism, Bayes' theorem can be used to show that Descriptive Humeanism is a posteriori much less likely to be true than Governing Non-Humeanism. Thus, regularities do more than just give rise to our concept of laws of nature; they can actually be used to demonstrate that certain theories of laws are a posteriori unlikely.

In Chapter 6, I complete the argument for Governing Non-Humeanism by comparing it to Descriptive Non-Humeanism. I begin by comparing the a priori probabilities of these theories. Next, I argue that observed regularities must be understood to be regularities among categorical properties (that is, among properties that are not bare dispositions). Finally, I argue that Descriptive Non-Humeanism has no explanatory power over such regularities; it is only able to explain regularities that hold among properties that *are* bare dispositions. Since Governing Non-Humeanism does provide a good explanation of regularities holding among categorical properties, it is the preferred view. Thus, if my methodology set out in Chapter 2 is correct, the final result is that Governing Non-Humeanism is the theory that is best supported by our empirical observations.

In Chapter 7, I provide a summary of the overall argument for Governing Non-Humeanism and discuss some metaontological implications of my empirical method of practicing metaphysics.

Chapter 2

Methodology

2.1 Introduction

In this chapter, I defend an empirical method of practicing metaphysics. The method avoids appeals to intuitions, ordinary beliefs, scientific theories, and the like. It does not accept basic principles of simplicity, unity, and the like. Instead, it proceeds from logic, analytic principles, and immediate experience alone. It shows that, contrary to common belief, (something like) British empiricism allows for the practice of a robust metaphysics. Though the primary aim of the method is to show how a traditionally empirical metaphysics can be vindicated, the method should be of interest to all metaphysicians. The core of the method demonstrates that immediate experiences are relevant to metaphysical theories, and this core is consistent with many epistemological positions. Thus a very interesting conclusion is preserved even if many of the central assumptions of this project are rejected: that immediate experiences are relevant to metaphysical theories, and thus that metaphysics can be empirically informed without appealing to realistically interpreted scientific theories.

This chapter is organized as follows. Section 2.2 explains my epistemological assumptions, which are essentially those of British empiricism. Section 2.3 discusses epistemic challenges to the practice of metaphysics unique to this version of empiricism. Most importantly, such empiricists must demonstrate how we can form metaphysically interesting concepts and how appearances (the things given in experience) constitute evidence for metaphysical theories (theories about the world itself). Section 2.4 introduces the tools and methods required for my empirical method.

Section 2.5 synthesizes these resources and demonstrates how the resulting method defeats the reasons for skepticism concerning genuine metaphysics. The central idea is that there are analytic truths of the form *if the world is in observable state O then there is a certain probability that metaphysical theory W , a theory about the world itself, is true*. We can know these conditionals a priori. Then, for certain O , we can confirm that the world is in state O and infer that there is a certain probability that theory W is true. In the Appendix to this chapter, I explain my stance on the methodological relevance of intuitions and ontological economy.

2.2 Epistemological Assumptions

The purpose of this section is to introduce the epistemological assumptions of this project. In the next section I shall explain some challenges to the practice of metaphysics specific to the epistemological position introduced in the present section.

I assume a robust analytic/synthetic distinction according to which a sentence is *analytic* if and only if it is true (false) in virtue of its meaning or logical form, and *synthetic* otherwise.

I assume a *foundationalist epistemology* according to which all inferentially justified beliefs are justified on the basis of non-inferentially justified beliefs. The nature of non-inferentially justified beliefs is explained by the assumptions that follow.

I assume two different empiricist theses. The first, *concept empiricism*, holds that all descriptive (that is, non-logical) concepts are either given in immediate experience or analyzable in terms of concepts given in immediate experience. The second, *doxastic empiricism*, holds that, ultimately, all justification of synthetic propositions is on the basis of immediate experience.¹ Thus all non-inferentially justified beliefs concerning synthetic propositions must be justified on the basis of immediate experience. But what account is given of immediate experience?

I assume that the objects of immediate experience are *sense data* (or *qualia*, though for our purposes it doesn't matter which). This assumption is important for two reasons. First, it provides

¹ Note that the latter thesis may need to be tempered slightly to fit with the other assumptions—most importantly, the assumption that we have justification for the truth of certain logical principles. However, there are good dialectical reasons for initially using this thesis in its strongest form.

a suitable foundation for both concept empiricism and doxastic empiricism. Given that the objects of immediate experience are sense data, it follows both that we can form clear and distinct concepts of them and that beliefs about them are non-inferentially justified. Not surprisingly, then, it is natural to adopt an *internalist* account of the language of metaphysical inquiry. Just as we stand in a privileged epistemic relation to our immediate experiences, so we shall stand in a privileged epistemic relation to the semantic content of the terms used to describe those experiences. This places us in the position to analyze the different terms in our language (or, as I shall sometimes speak, the different concepts we possess) for analytical connections. Second, this assumption entails that there is a sharp distinction between the objects of perception and the external world. In other words, it drives a wedge between *appearance* and *reality*.

Finally, I assume that we have good reason to believe that some particular system of logic is correct (it does not matter which, as long as the system provides sufficient resources for the technical apparatuses introduced later), or at least that we have good reason to accept certain basic rules of inference. This assumption is indispensable. Among other reasons, it is required for the justification of inferentially justified beliefs. One might worry that it undermines the claim to be offering an empirical method, since it can be argued that the matter of choosing the correct system of logic requires synthetic a priori reasoning (and thus that the assumption of doxastic empiricism requires modification). Perhaps this is true. But it does not reflect the practices of the many empiricists who appeal to logical principles (in support of epistemic reasons) in their various levels of inquiry. And, even if it were true, a moderate empiricism that endorses the synthetic a priori only for such matters could still proceed in accordance with the method without introducing any additional synthetic a priori justification.

For later reference, let *empirical foundationalism* be the position that accepts the above assumptions. This is essentially the view of the early-modern British empiricists, Locke, Berkeley, and Hume.² (I decided not to call my position *British empiricism* because I do not intend to be reconstructing or interpreting the views of the British empiricists. That said, I do believe that my

² It is also, more or less, Bertrand Russell's view; see for example (Russell 1912, chapter 5).

position is essentially the same as theirs.) Since Hume, many philosophers have thought that this epistemic position precludes the practice of metaphysics. In the next section I shall explain why.

Before continuing, I want to make an important disclaimer. In assuming empirical foundationalism, I am making some very strong epistemological assumptions. However, the core of the method I shall defend does not in fact require empirical foundationalism. Thus, although empirical foundationalism deserves careful scrutiny, this is not the place for such scrutiny. Suppose we weakened some of the assumptions. Suppose we rejected concept empiricism, doxastic empiricism, and the analytic/synthetic distinction. The method would not then show that a *purely empirical* metaphysics is possible, but it would show that metaphysics can be *empirically informed* in a special way; namely, it would show that immediate experiences are relevant to metaphysical theories. Thus the method I shall defend should be of interest even to those metaphysicians who do not share my epistemological assumptions.

2.3 An Argument Against Empirical Metaphysics

This section quickly sketches a familiar Humean (“Humean” in the sense that it is due to Hume, not in the sense that it presupposes Humeanism) argument against the possibility of metaphysics that is based on the distinction between appearance and reality accepted by empirical foundationalism.

It is difficult to say exactly what metaphysics is, but for our purposes the following characterization shall suffice. *Metaphysics* is the area of study concerned with questions such as the following: does God exist? is there cause and effect in the world? what is the nature of cause and effect? are there properties? what is the nature of properties? are there minds? what is the nature of mind? are there features of the world that answer to modal concepts? what is their nature? These are questions that appear to resist scientific answers, and so the empiricist must have a different empirical method in order to answer them.³

³ This is not to say that scientific theories are irrelevant to metaphysics. However, my view is that scientific theories—that is, scientific theories with all their philosophical baggage (or adornments) removed—are essentially just systematizations of observation statements. In order for such a systematization to become relevant to metaphysics,

On the surface, the paradigm metaphysical questions listed above are questions not about mere appearances (they aren't questions about sense data or qualia), but rather about reality, the world itself. According to my assumptions, we do not immediately observe the the world itself, much less its metaphysical entities. At the very least, we do not discern the nature of supernatural phenomena, causal relations,⁴ time, modal relations between contingent entities, and the like through immediate observation. Regardless, the upshot is that matters metaphysical are not resolved through immediate observation. The questions of metaphysics just aren't about that which can be given in immediate experience. Whatever account is given of appearances, it does not entirely capture the subject matter of metaphysics, because there is an epistemic gap between appearance (the foundation for all our descriptive concepts) and reality. Accordingly, one problem for empirical foundationalism is that it needs to provide an account of theoretical concepts that is consistent with concept empiricism; otherwise, we can't even talk about reality, much less have justification for beliefs about it.⁵

(Note: For this reason, many philosophers have thought that empiricists who practice metaphysics are merely engaged in conceptual analysis. If we like, we can distinguish *genuine metaphysics* from *metaphysics as conceptual analysis* as follows: genuine metaphysics makes claims about the world itself, for example, God himself, causes themselves, etc.; metaphysics as conceptual analysis does not, restricting itself to claims about our concepts of God, causes, etc. There is a gap between our epistemological foundation—the non-inferentially justified beliefs about our experiences—and the world itself. Genuine metaphysics requires the gap to be bridged; metaphysics as conceptual analysis does not.)

Another, perhaps more serious, problem is that many philosophers have thought that bridging the gap between appearance and reality is impossible, at least for the empirical foundation-

it must first be given a philosophical interpretation. However, interpreting scientific theories in a way that is relevant to metaphysics is very close to metaphysics itself, if it is not the very same thing.

⁴ *Pace* Anscombe (1971), Armstrong (1983), and Fales (1990, Chapter 1). My opinion is that these philosophers have not succeeded in countering Hume's well-known argument to the effect that causal relations are not observable (found in Section IV of the *Enquiry*). More on this later.

⁵ This would be so even if Berkeleyan idealism were true. For in that case, one could not recognize the objectivity of appearances without theoretical concepts, nor could one give an account of the theory itself: namely, that appearance *is* reality.

alist. That is, they have thought that appearances cannot justify beliefs about reality—that appearances radically underdetermine the nature of reality. The basic reasons may be expressed as follows. On one hand, it seems that analytic principles aren't suited to the task at hand. Analytic propositions don't entail synthetic propositions, and 'the gap between appearance and reality is bridged' expresses a synthetic proposition. On the other hand, all synthetic principles are a posteriori, but these principles are concerned only with appearances, not with the world itself. Being concerned solely with appearances, they cannot bridge the gap. Therefore, it looks like any principle that *could* bridge the gap would have to be a synthetic a priori principle. Since empirical foundationalism precludes justified beliefs concerning such principles, it looks like empirical foundationalism precludes justified beliefs concerning genuine metaphysics.⁶ Thus we have a serious challenge for doxastic empiricism even if the problems for concept empiricism can be solved.

A demonstration that the above argument is mistaken requires a *linking principle* that connects the world of appearances to reality, the world itself. I shall argue that there are analytic linking principles of the form: $O \supset W$, where O is an empirical observation statement and W is some proposition about the (probable) truth of a theory about the world itself.⁷ It is the conjunction of this sort of principle with actual empirical observations that provides the bridge between appearance and reality. This strategy shouldn't come as a surprise. The argument above is sound only if it is assumed that metaphysical theories cannot have important implications for our experience. However, the analytic linking principles are supposed to show precisely the opposite: that some metaphysical theories *do* have important implications for our experience—that is, that some metaphysical theories are *not* underdetermined by appearances.

⁶ This type of argument is found in Carnap's famous "The Elimination of Metaphysics Through Logical Analysis of Language" (1959, 76).

⁷ Strictly speaking, the linking principles will be atomic statements of conditional probability, so the precise form is as follows: $P(T|O)$, the probability of theory T given observation(s) O .

2.4 Empirical Resources

The purpose of this section is to introduce some resources available to the empirical foundationalist. I cannot offer anything like a complete defense of the thesis that these methods are consistent with empirical foundationalism, but I will provide at least some justification of this claim.⁸ Along the way, I shall explain how these resources can accommodate concerns that arise for the metaphysics of laws of nature.

2.4.1 Empirical Foundationalism

The **first method** provides a way of forming basic concepts; it is nothing more than the parts of empirical foundationalism that state the criteria for the formation of basic concepts—namely, its empiricist and foundationalist components. Following this method ensures that we stand in privileged epistemic relations to our (properly formed, basic) concepts. In short, empirical foundationalism provides us with a solid epistemic foundation for those concepts having to do with the world of appearances.

The **second method** provides a way of forming complex concepts; it involves the system of analysis that is used to generate new concepts and justify inferentially justified beliefs. For our purposes, this method has two distinct components: first, a system of logic that permits quantification over properties in addition to individuals;⁹ second, a method of defining *quasi-logical concepts* such as *parthood*, *property*, *event*, *state*, and *particular*.¹⁰ Regarding the second component, empirical definitions of the terms mentioned is possible. For example, we could give the following definitions: the concept of *parthood* is defined either by a special sort of ostension (e.g., by distinguishing features of one's visual field) or using a modal definition such as x is a part of y $=_{df}$ it is impossible that y exists as it presently is and x does not exist; x is a *particular* if and only if it is logically possible for there to be a y which does not differ qualitatively from x , but which

⁸ I hope to provide a more complete defense of this claim in my next major research project.

⁹ Two points: First, I assume that the adoption of this logic does not preclude nominalism; one can use such a logic without reifying properties. Second, it is worth noting that Hume did not have this method at his disposal.

¹⁰ This list is borrowed from (Tooley 1987).

is not identical with x ; x is a (*universal*) *property* if and only if it is logically impossible for there to be a y which exactly resembles x , but which is not identical with x .¹¹ These two components allow for the analysis of a great many concepts; they even allow for concepts of entities that are in principle unobservable, such as the concept *parts too small to be observed*. More importantly, they provide the right sort of foundation for the third method (discussed in 2.4.2).

2.4.2 Laws as Theoretical Entities

David Hume thought that we could not learn anything about laws of nature as they exist in the world. The basic reason is that we do not directly observe laws or causes; we merely observe regularities. This section begins with an explanation of the need for a theoretical treatment of nomological concepts. This is to say that Hume was *right* to think that our causal and nomological concepts are in need of analysis. This section concludes with an explanation and defense of a version of the Ramsey/Carnap/Lewis method of defining theoretical terms that is consistent with a realist interpretation of theoretical entities. This is to say that Hume was *wrong* to think that the only possible analysis of nomological concepts entails Descriptive Humeanism, a more general version of Hume's well-known regularity theory of laws.

2.4.2.1 Nomological Concepts: The Need for Analysis

Empirical foundationalism entails the following view of complex concepts: all complex concepts must ultimately be grounded in concepts that are analytically basic, that is, not in need of analysis. (We could elaborate on the idea of what it is for a concept to be analytically basic by saying that a concept is analytically basic if and only if analysis of that concept is both impossible and unnecessary.) According to empirical foundationalism, the only descriptive analytically basic concepts are those that arise on the basis of immediate sense experience; immediate sense experience involves being in direct perceptual or introspective contact with the object or property in

¹¹ The list of quasi-logical concepts is borrowed from (Tooley 1987), the latter two definitions are borrowed from Tooley's introduction to (Tooley 1999, Volume 3), and the latter definition of 'parthood' was suggested by Tooley in correspondence. The point here is not to defend these definitions, but rather to show that it is possible to give empirically respectable definitions of quasi-logical concepts.

question.¹² For example, consider how it is that we come to understand color concepts. We learn about color through our immediate experiences of color. Words that refer to colors are defined via ostension. For example, we see a set of red objects and are told that that quality which our experiences of them have in common is the referent of 'red'. Without the relevant experiences, this definition would be worthless.

We can now discuss nomological concepts. Are any of them analytically basic? Some have thought so. For example, Anscombe (1971), Armstrong (1983), and Fales (1990, Chapter 1) have argued that we directly experience singular causal relations—a species of nomic relations—and that these experiences form the basis for all further nomological concepts. We may call this view of nomological concepts *direct realism*. The appeal of the direct realist approach is that it provides a solid epistemological foundation for all nomological concepts. If direct realism is true, some nomological concepts are analytically basic, and these form the foundation for more complex nomological concepts. If direct realism is true, (at least some) nomic entities are *not* theoretical entities; we believe in them because we experience them directly, not solely because they are postulates of a theory that we accept. Direct realism solves many of the epistemological problems often associated with Non-Humeanism and Governing Humeanism. Unfortunately, I believe that direct realism is false. I shall argue that no nomological concept (of ours, anyway) is analytically basic. The result is that, if any version of Non-Humeanism or Governing Humeanism is to be entertained at all, a theoretical treatment of nomological concepts is required. If direct realism is false, then Non-Humean or Governing Humean nomic concepts must be theoretical concepts (complex concepts which require a special sort of analysis).

The argument against direct realism explained here is an ancestor of Hume's. The particular version is due to Michael Tooley (Forthcoming, Chapter 2). Tooley begins by pointing out that it is not sufficient to say that our normal sense experience *gives rise to* the concept of causal relations. To establish direct realism, it must be shown that this concept *arises immediately*, and not on the

¹² "If a concept is analytically basic, then, by definition, one can acquire the concept in question only by being in perceptual or introspective contact with an instance of the property or relation that is picked out by the concept" (Tooley Forthcoming, Chapter 2).

basis of some inference. There is thus a crucial distinction between

(1) concepts arising through sense experience

and

(2) concepts that are analytically basic.

Though analytically basic concepts arise directly from sense experience, it is not the case that all concepts arising through sense experience are analytically basic. Hume famously thought that our beliefs about cause and effect were precisely of the sort that “arise” through experience but that are not given in experience (see *Enquiry* Section IV). We don’t *perceive* any connection between cause and effect; we just *imagine* it to be there. The concept of cause and effect arises through a combination of sense experience and imagination, but it is not, according to Hume, analytically basic. Fales’s (1990) favorite example of a purportedly analytically basic nomological concept involves the sensation of tactile pressure. The basic idea is that the experience of tactile pressure naturally gives rise to the ideas of cause and effect because the causal relation itself is observed directly. Of course, we must be careful not to conflate (1) and (2). The fact that the concept of singular causation arises through our experience of tactile pressure does not on its own establish that it is analytically basic.

Tooley then argues that once one is aware of this distinction, it becomes apparent that no nomological concept is analytically basic. The argument is easier to understand if we start with an analogous case. Consider the concept of a physical object’s being red: call this concept *physical redness*. This concept could be acquired even if there were no red things in the world—perhaps due to hallucinations or deceptions by an evil demon—and hence the concept of a physical object’s being red is not analytically basic since it can be acquired without perceptual or introspective contact with an instance of the property of physical redness. This is in stark contrast to the concept of redness as a quality of experiences, which is analytically basic; call this concept *qualitative redness*. Thus, qualitative redness and physical redness are distinct concepts since only the former is analytically basic. This result is further supported by contemporary physical theories of color,

according to which the color of an object is a complex matter having to do with properties of its microphysical structure. Thus the concept of physical redness is complex. This is not the case for qualitative redness, so once again we have good reason for thinking that physical redness and qualitative redness are distinct.

Now consider the direct realist's claim that the experience of tactile pressure on one's skin is a case in which the singular causal relation is directly experienced, and is thus analytically basic. Here we have two concepts: the concept of *tactile pressure* and the concept of *singular causal relations*. Suppose that you experience tactile pressure as you press your hands together. Have you directly experienced a singular causal relation, or have you merely experienced the tactile pressure? It is quite easy to imagine circumstances in which one would *feel* as though tactile pressure was the cause of some indentation in one's skin even though no such causal relation had occurred: one might be hallucinating, in which case no causal relation underlies the experience. Though one cannot be deceived about whether one is experiencing the tactile pressure, this does not show that the concept of singular causation as a relation between objects *in the world* is analytically basic; it merely tells us something about the quality of our experience of tactile pressure.¹³ Therefore, the concept of tactile pressure does not serve as an adequate base of nomological concepts, just as the concept of qualitative redness doesn't serve as an adequate base for the concept of physical redness.

It seems likely that the preceding argument will apply to every (realistically construed) nomic entity in addition to singular causal relations. None of the following appear to be good candidates for directly observable entities: governing laws of nature, bare dispositions, counterfactuals, et cetera. The entities that satisfy these concepts are, so far as we can tell, unobservable. Thus, the preceding line of argument would apply to any of these concepts. For each one, there is a Cartesian scenario in which one can have a non-veridical experience of the relevant entity (that is, the relevant "causal relation", "law", et cetera);¹⁴ hence, no nomological concepts are

¹³ Tooley develops this argument with greater precision. Here, I have merely explained its basic structure.

¹⁴ Strictly speaking, this is a bit imprecise, since we do not ever experience these entities. Rather, we merely imagine that we are experiencing such entities.

analytically basic.

The result is that direct realism is false; nomological concepts are not analytically basic. Therefore, they must be analyzed in terms of concepts that are more basic. According to Descriptive Humeanism, the truthmakers for nomological propositions concerning observed regularities ultimately reduce to those given immediately in sense experience; nomological propositions that concern observed regularities do not require theoretical treatment. Hume's regularity theory of laws is a well-known example; laws are nothing (or little) more than generalizations of observation statements. For the Non-Humean and Governing Humean, however, there is a further requirement. In light of the above argument, no entity uniquely posited by the Non-Humean or Governing Humean can be immediately observable. Therefore, such theorists require a method of explaining our concepts of such entities. Fortunately, a theoretical treatment of the relevant concepts is available that allows us to understand the concepts and use them to refer to the objects that satisfy them. Thus, Non-Humeanism and Governing Humeanism do not require direct realism in the first place.

2.4.2.2 Defining Theoretical Terms

The Ramsey/Carnap/Lewis method of defining theoretical terms allows for empirically respectable analyses of our concepts of unobservable entities. See (Ramsey 1931), (Carnap 1966), and (Lewis 1970); Tooley (1987, Chapter 1) provides a nice summary of objections to this method along with replies. The importance of this method cannot be understated. It allows the empirical foundationalist to form concepts of (and refer to) unobservable entities, such as the metaphysical entities in the world itself. Without this type of method it would not be possible to provide empirically respectable definitions of Non-Humeanism or Governing Humeanism. This method of defining theoretical terms is our **third method**.

Before explaining the method, a clarification will be helpful. A *theoretical entity* is some entity that we believe in because it is the referent of a term in a theory we accept (see (Lewis 1970, 428)). On the other hand, a *theoretical term* is one whose definition is special in that its meaning

depends on the theory in which it appears; such terms are introduced into the language through their inclusion in a theory. For example, the term 'H₂O' had no meaning prior to its inclusion in molecular theory, though 'water' certainly did. There are thus two distinct questions one can ask: Are there any theoretical entities? and How do we understand theoretical terms? The answer to the first depends on whether there is any true theory; if we are justified in believing that a theory is true, we are justified in believing that there are theoretical entities of the sort specified by that theory. The answer to the second question is specified by the method, its goal being to provide definitions of theoretical terms using only those terms with which we are already familiar. So, which terms are already familiar? Those specified by the first two methods above. Lewis's (1970) name for this set of concepts is the *O*-terms: *ordinary* terms, *old* terms, *original* terms. We shall not restrict ourselves to any old terms, but only those that are empirically respectable in accordance with the first two methods.

The precise method of defining theoretical terms is somewhat technical, though it is based on rather simple ideas. After explaining the method itself, I shall discuss a method that ensures that the terms used in the construction of a theory are definable within the Ramsey/Carnap/Lewis method of defining theoretical terms. The result is that there are relatively few occasions in which the technical method itself must be directly applied, and thus the technical elements can be avoided in usual practice.

What follows is a brief and (inasmuch as possible) non-technical sketch of Ramsey's version of the method. The first step is to treat all theoretical terms in a theory uniformly by paraphrasing all statements including theoretical terms into the same logical/syntactical form. Consider a theory *T* which ordinarily includes (3). The theory could have included (4) instead, since (3) and (4) are equivalent.

(3) *a* is an electron.

(4) *a* has the property *electronhood*.

The purpose of this step is merely to simplify the method. It allows for the uniform syntactical

treatment of all theoretical terms.

The second step is to take the conjunction of the statements of T that include theoretical terms (the purpose of doing so is simply to put all of the theoretical terms in one place for ease of reference). This can be represented as follows, where each τ is a theoretical term:

The Postulate of T : $T[\tau_1, \dots, \tau_n]$

The Postulate of T can be existentially quantified to eliminate theoretical terms. This yields

The Ramsey Sentence of T : $\exists x_1 \dots \exists x_n (T[x_1 \dots x_n])$

The basic proposal, due to Ramsey (1931), is that the Ramsey Sentence of T may be used in place of T itself—that is, in place of the Postulate of T . No one has any trouble understanding the Ramsey Sentence of T ; it does not include any theoretical terms. The result? “[I]nsofar as the theory T serves as a device for systematizing O -sentences, the Ramsey sentence of T will do the job as well as the postulate itself.” (Lewis 1970, 431)¹⁵

Ramsey/Carnap/Lewis theoretical terms succeed in referring to theoretical entities (if there are theoretical entities) by way of descriptions. The descriptions are special in that they do not refer in virtue of the intrinsic properties of theoretical entities, but in virtue of their *extrinsic* or *relational* properties to objects or concepts toward which we stand in a privileged (and empirically respectable) epistemic relation. The descriptions are entirely relational. They pick out theoretical entities in virtue of the fact that those entities instantiate relational properties analyzed in terms of concepts that are empirically respectable.¹⁶ These features of the method have a very important implication. Theoretical terms are semantically equivalent to descriptions, and, in many cases, the semantic content of these descriptions is *synthetic* (this is to say that a proposition attributing the relevant description to an object is a synthetic proposition). The result is as follows: *different theories may have different implications for our experience*. This is very important; a synthetic definite

¹⁵ Lewis’s extension of Ramsey’s method, which I prefer, is slightly more precise and thus gives a more complete understanding of theoretical terms, but it is rather complicated. For our purposes, however, Ramsey’s version of the method will suffice.

¹⁶ Michael Tooley has helpfully pointed out that some theoretical terms can be analyzed in terms of the relations among the relevant theoretical entities to each other, rather than in terms of their relations to empirically respectable things.

description stated in the *O*-language will have observable consequences. This allows for the possibility that distinct metaphysical theories vary in their empirically observable consequences, since distinct metaphysical theories will be (in the right cases) equivalent to distinct synthetic descriptions. Therefore, the method of defining theoretical terms ultimately allows for the possibility of the empirical confirmation (or falsification) of paradigmatically metaphysical theories.

The basic ideas underlying the method are not as technical as they might appear. The method simply uses a special sort of description—one that even the empiricist can understand—in order to pick out theoretical entities.¹⁷ Once we are aware of the method we can bypass it by constructing careful definitions within the *O*-language. For example, suppose that we define an atomic Non-Humean property as follows:¹⁸

AN: *X* is an *atomic Non-Humean property* if and only if *X* is a property such that its instantiation at a spacetime point entails the existence of a property instantiation at a distinct spacetime point.

and added that

S: there is at least one atomic Non-Humean property.

(Since our concern is with Non-Humeanism, let's assume that the concept of a spacetime point is part of our *O*-language.) We aren't in the position to observe entailment relations between logically distinct states of affairs; hence, 'atomic Non-Humean property' refers to a theoretical entity if it refers at all. However, *AN* does not appeal to any term in the *T*-language. The only theoretical term in the definition is the definiendum—'atomic Non-Humean property'—itself (it is important to remember that logical and quasi-logical vocabulary is part of the *O*-language). *S* uses the term 'atomic Non-Humean property', but otherwise it merely makes an existential generalization. Together with *AN*, *S* seems to make perfectly good sense. So why do we need the method at all?

¹⁷ The Ramsey method allows for theories to be multiply realized. In effect, it is as though the descriptions are *indefinite*. Lewis's method fixes this problem; the resulting descriptions are, in effect, *definite*.

¹⁸ This is not the correct way to define atomic Non-Humean properties—it is too narrow—but this definition makes for a simple example.

Let's consider an application of the Ramsey/Carnap/Lewis method. Suppose that T includes the following postulate:

- (5) Point i instantiates N , and this entails an observable state of affairs O such that $Ni \neq O$; that is, $Ni \& (Ni \supset O) \& Ni \neq O$.

N is the theoretical term which refers to an atomic Non-Humean property; i and O are terms in the O -language. We can Ramsify this postulate, yielding

- (6) $\exists X(Xi \& (Xi \supset O) \& Xi \neq O)$.

Note that (6) is, effectively, just saying that the definition of an atomic Non-Humean property is satisfied. But of course, we can say this in our ordinary language without explicitly appealing to the process of Ramsification; this is precisely what the conjunction of AN and S tells us. Stating the definition in the O -language is quite helpful; it does everything that Ramsification does. Thus, practically speaking, we can often bypass the Ramsey/Carnap/Lewis method. If we go about our metaphysics in the right way, taking care to precisely define all concepts using only terms from the O -language (as the good empiricist should be doing anyway), then we won't have need to apply the technical method itself. This is not to say that the method isn't important; in fact, the method helps us to set the criteria for forming suitable definitions of theoretical entities using only terms in the O -language. After all, this is effectively what the process of Ramsification accomplishes (as does Lewis's improved version of the method, with a greater degree of precision).

Thus far, I have assumed that the purpose of the method just explained is to explain how a *realistic* construal of theoretical terms is possible. However, it is important to note that the method itself is not prejudiced against anti-realist or reductionist construals of the referents of theoretical terms. Nothing in the method itself guarantees that theoretical terms have a denotation; one could use the method to define (that is, explain how we understand) the terms without holding that the theory is *true*. If the theory isn't true, we aren't forced to say that the terms denote anything. Thus anti-realism is consistent with the method. Furthermore, nothing in the method itself guarantees that there is no analysis or reduction of theoretical *entities* in terms of the properties and relations

of observable objects; thus, reductionism is consistent with the method.¹⁹ This is not to say that all theories generated by the method are consistent with Descriptive Humeanism. Some are not. I believe that Hume was right to demand an analysis of nomological concepts in terms of concepts that are, ultimately, non-nomological. However, the empiricist's preference for analysis of all concepts in terms of those that are immediately given in sense experience entails neither Humean Anti-Realism nor Humean Reductionism nor Descriptive Humeanism. We have in hand a method of defining theoretical terms that is consistent with a realistic construal of theoretical entities.

2.4.3 The Method of Confirmation

We have a method that allows us both to understand our theories and to determine their empirical implications. But how do we use empirical observations to confirm one theory (or set of theories) over the others? I shall quickly set out and defend a Bayesian method of confirmation (this is our **fourth method** consistent with empirical foundationalism).²⁰ It might be thought that this method is so well known that an explication of it is unnecessary. However, there are certain features of the method that have been, from the perspective of philosophers of science, unsatisfactory with regard to the confirmation of scientific theories. Specifically, some philosophers of science (and many scientists) have adopted a falsificationist view of scientific theories according to which theories cannot be confirmed, but merely falsified (see (Popper 1959)). To a certain extent, the questions with which we shall be concerned avoid this criticism entirely, so it is worth explaining how the Bayesian method of confirmation may be applied to our subject in a way that will satisfy the falsificationist.

Popper (1959) believed that scientific hypotheses could be falsified, but that they could not be confirmed. That is, he believed that empirical investigation could provide evidence *against* a

¹⁹ Of course, the *content* of T might be such that, if T is true then reductionism is false. But this is a desirable consequence.

²⁰ By 'Bayesian' I just mean the method that involves updating probabilities in light of evidence in accordance with the theorems of the probability calculus. I do not mean to be advocating a *subjective interpretation of probability*; that matter is discussed in the next section.

hypothesis, but that it could not provide evidence *for* a hypothesis. The basic reason was that Popper thought that genuine confirmation required a solution to the problem of induction—a solution which we do not have. There are, however, circumstances in which confirmation is possible within a falsificationist framework even if one lacks a solution to the problem of induction. Consider, for example, the rule of *disjunctive syllogism* in classical logic. This rule says that, given a disjunction and the negation of one disjunct, one may infer the other disjunct. Let $H1$ and $H2$ be competing hypotheses, and suppose one knows that either $H1$ or $H2$ is true. Now suppose that one learns that $H2$ is false due to its incompatibility with empirical evidence. $H2$ has been falsified, $H1$ has not; this entails that $H1$ has been confirmed, and no appeal to induction is required. So says *disjunctive syllogism*. If our situation is like this—namely, that we know that the disjunction is true, and then we come to have evidence against one of the hypotheses—there is no problem of explaining, in general, how confirmation works. This sort of confirmation is acceptable to the falsificationist.

There are two major difficulties that must be met in order for one to employ the rule of disjunctive syllogism to confirm a theory. First, one must be justified in believing that the disjunction is true; unfortunately, this often requires the set of competing hypotheses to be exhaustive, and its members exclusive. This poses a serious problem for the scientist. The reason is that scientific theories are very complex. The degree of complexity and precision of a hypothesis is directly related to the number of competing hypotheses. For example, the hypothesis that the force of gravity is represented by Newton's formula

$$F = G \frac{m_1 m_2}{r^2}$$

is much more precise than the hypothesis that there is *some* force involving the mass of physical objects. There are infinitely many competitors of the same level of generality as the former hypothesis (G could be replaced by any number); there is only one competing hypothesis of the same level of generality for the latter: the hypothesis that there is *not* any force involving the mass of physical objects. The more general hypothesis is much simpler, and much easier to confirm,

than the more specific hypothesis. Unfortunately, the more general hypothesis is virtually useless. Scientists won't care much about it at all. But the metaphysician might. The metaphysician is interested in a different sort of question entirely. She may not care whether her theory ever leads to successful predictions. That's not the (primary) purpose of metaphysics. She is interested in whether her theory is *true*.²¹ It may be the case that we metaphysicians must settle for very general metaphysical hypotheses. But, of course, most of the hypotheses we are interested in are quite general: Is naturalism true or false? Does God exist or does he not? Do we have free will or do we not? Is causation a real feature of the world or do we merely imagine it? I would count the discipline of metaphysics a success if it could provide answers to questions such as these.

Fortunately, solving this difficulty is a rather simple matter for the metaphysician; she can just define her theories in such a way to guarantee that this condition is met. The metaphysician defines the relevant hypotheses so as to form a set of mutually exclusive and jointly exhaustive hypotheses. Thus, the disjunctive syllogism schema can be used to genuinely confirm metaphysical hypothesis without a solution to the problem of induction.

The second difficulty is that hypotheses are rarely falsified in the strict sense explained above. It is rare to find empirical evidence that flat-out contradicts the implications of a hypothesis. Such a contradiction usually is derivable only with questionable auxiliary assumptions. (Note that this will not be an issue if the theories compared are mutually exclusive and jointly exhaustive.) Thus, it is generally thought that evidence is *probabilistic*. For example, we think it merely *unlikely* that Newtonian physics is true in light of the Eddington experiments, not *impossible*.

Fortunately, there is another formal system that can be applied to our purposes. This is the probability calculus.²² The theorems of this system are no weaker, no less reliable than those of classical logic; though they concern probabilities, they do not rely on a solution to the problem of induction. For our purposes the most important theorem of the calculus is Bayes' theorem, which

²¹ Of course, the scientist may be interested in this same question. But it is important to note that the scientists is *also* interested in the predictive power of her theory, whereas the metaphysician is not.

²² See (Howson and Urbach 2006, Chapter 2) for an accessible introduction to the probability calculus.

is often stated in the following form:

$$P(p|e) = \frac{P(p)P(e|p)}{P(p)P(e|p) + P(\sim p)P(e|\sim p)}$$

Bayes' theorem is important because it describes the degree of support that evidence e has for a given hypothesis p . Having solved the first difficulty for confirmation, Bayes' theorem can be directly applied to metaphysical hypotheses since the questions posed will concern mutually exclusive and jointly exhaustive hypotheses. This is a crucial point, since, once again, the application of Bayes' theorem does *not* presuppose a solution to the problem of induction. The basic idea here is this. Though sometimes evidence against a theory is not strong enough to falsify that theory, it can nevertheless make the theory less likely than it would have been in light of different evidence, or no evidence. If the theories are mutually exclusive and jointly exhaustive, evidence against one theory—even if it is merely probabilistic evidence—constitutes evidence in favor of the other. So says Bayes' theorem. Thus, the general spirit of the disjunctive syllogism schema is preserved in probabilistic cases. (Again, the metaphysician has an advantage over the scientist here, since she can just define her theories so as to be mutually exclusive and jointly exhaustive.)

Of course, we can't apply Bayes' theorem to determine whether an observation is evidence in favor of one hypothesis or another unless we know the degree to which the competing hypotheses *predict* that evidence—that is, unless we know the conditional probabilities on the right-hand side of Bayes' theorem; furthermore, we can't apply Bayes' theorem to determine the probability of a hypothesis given the evidence unless we know the prior probabilities of the hypotheses in question. The theorems of the probability calculus are useless to our purposes unless we have a method of assigning the relevant probabilities in the first place. Such a method is called an *interpretation of probability*. In the next section, I explain and defend the method I prefer.

2.4.4 Assigning Probabilities

There is a fair amount of controversy as to how probability should be interpreted. I shall begin by explaining the relevant concept of probability as a sort of *epistemic probability*. I then

set out my preferred approach to assigning epistemic probability: a version of the *logical interpretation* of probability set forth by Carnap (1962). Then I shall discuss some objections to the logical interpretation, discuss alternatives, and explain why the general approach that underlies the logical interpretation is the most satisfactory.

2.4.4.1 Preliminaries

There are different ways to understand statements about probability. My explanations are based on those given by Mellor (2005, Chapter 1). First, statements of probability may be understood as expressing degrees of belief, or *epistemic credences*. An epistemic credence tells us how strongly one believes a proposition. This is the best way to interpret statements involving betting odds: a rational person is willing to place a bet at certain odds if and only if one's degree of belief differs from the odds in question. For example, if a bookie tells you that the odds are 1:1 that *X* will win, you will place a bet on *X* only if you believe that there is a greater than 50% chance that *X* will win.

Second, statements of probability may be understood as expressing the (objective) *epistemic probability* of a hypothesis in light of evidence (or lack of evidence). Epistemic probabilities “measure how far evidence confirms or disconfirms hypotheses about the world... But they are not mere matters of opinion: whether, and to what extent, evidence counts for or against a hypothesis looks like an objective matter” (Mellor 2005, 8). To say that an epistemic probability is *objective* is to say that two (equally imaginative and equally intelligent) individuals with all the same background information and same evidence would not be justified in assigning different probabilities to the proposition in question. Within the context of empirical foundationalism, it is easy enough to explain why these statements are objective: they will ultimately be grounded in one's language that is generated in an empirically respectable way (but more on this later). For example, propositions about the likelihood that God exists given the fine-tuning of the universe are plausibly interpreted to be statements of epistemic probability.

Third, statements about probability may be understood as expressing *objective chances*, or

real probabilities that hold in the world. For example, the probability that a coin will land heads is usually thought to express an objective chance. There are cases in which these three forms of probability come together. In the case of a fully-informed ideal epistemic agent, they should coincide. But since we are not ideal epistemic agents, it will be best to treat them separately.

I shall be concerned primarily with epistemic probabilities. The reason is that our ultimate concern is with metaphysical statements such as the hypothesis that Non-Humeanism is true or that the probability of Humeanism given natural regularities is low. I shall avoid epistemic credences for the simple reason that we are trying to determine what our epistemic credences *should be*. To begin with epistemic credences would be to put the cart before the horse. I shall avoid objective chances for two reasons. First, it is not obvious (at this point) that we have epistemic access to the objective chances of the world itself. Second, many philosophers have thought that many of the metaphysical hypotheses under consideration are either necessarily true or necessarily false; this is to say that their objective chance of obtaining is either 1 or 0. But this will be of little help. For example, Goldbach's Conjecture has an objective chance of either 1 or 0 (since it is either necessarily true or necessarily false). But for our purposes, it will be best to treat such statements as expressing epistemic probabilities, such as the likelihood from our epistemic perspective that Goldbach's conjecture is true (note that this is actually a disguised conditional probability; hence, it fits the definition of epistemic probability given above). This allows for the full range of probabilities (greater than or equal to 0, less than or equal to 1); we can assign a number other than 0 or 1 to this hypothesis. In this sense, one can say, for example, that the epistemic probability that Goldbach's Conjecture is true equals .5. In my opinion, this is the only satisfactory way to treat metaphysical hypotheses, at least if one is interested in the project of evaluating these hypotheses with respect to their fit with evidence. (However, it will turn out that, if my preferred method of practicing metaphysics is correct, certain epistemic probabilities match up with objective chances.)

2.4.4.2 The Logical Interpretation of Probability

I shall utilize a *logical interpretation* of probability to assign epistemic probabilities. This is our **fifth method**. According to the logical interpretation, objective epistemic probabilities may be assigned through logical and semantical analysis, so it is an (analytic) a priori method. This subsection explains the basic features of this interpretation—modeled after Carnap’s (1962) interpretation—and shows how it may be developed to accommodate the debate between our four theories of laws explained in Chapter 1.

The logical interpretation of probability is based on the idea that one can determine the epistemic probabilities of competing theories by comparing the number (and perhaps kind) of epistemic possibilities countenanced by the competing theories. How can an empirical foundationalist give an account of the concept of an epistemic possibility? Carnap’s basic proposal, which I shall adopt, is that epistemic possibilities are determined by logical and semantical analysis. We are given a set of basic individual terms Σ and basic predicate terms Φ . We stipulate that these terms are those formed in accordance with the principles of empirical foundationalism. Carnap calls a maximal consistent permutation of the members of Σ and Φ (with the members of the former instantiating the members of the latter in the familiar way) a *state description*. State descriptions are our epistemic possibilities.

Put another way, a *state description* is a proposition describing an epistemically possible state of affairs or a set of epistemically possible states of affairs. (As a heuristic, we can think of a *state of affairs* as being an instantiation of a property or relation by some thing or ordered tuple of things,²³ but this definition is not intended to impart any ontological commitment.) Our interpretation of probability must tell us two things: how to determine what state descriptions there are, and how to assign probabilities to these state descriptions.

In one sense, settling the first issue—the matter of determining which state descriptions there are—appears simple: since we are concerned with objective epistemic probability, the state

²³ The neutral term ‘thing’ is used instead of ‘individual’ so as not to rule out higher-order states of affairs.

descriptions are just the epistemically possible ones. However, we need to say a little bit more about what constitutes an epistemically possible state description by saying more about the parameters Σ and Φ . Carnap assumes that the predicates in Φ must be monadic, that only individuals of order 0 satisfy predicates, and that all predicates are of order 1. I do not adopt these assumptions. The result is a language that has more descriptive power and is more complex than Carnap's. Further, we shall ultimately be interested in the general question How many state descriptions are there, *period*? To answer this question, the two parameters must be allowed to "float" over the range of possible values. Thus the quantity of state descriptions (that is, epistemic possibilities) is a function of a *range* of epistemically possible values for Σ and Φ .

This is all well and good, but as explained it is not clear how the system will allow us to describe Non-Humean state descriptions—that is, to describe all of the epistemic possibilities relevant to the debate concerning laws of nature. We can remedy this problem with two simple augmentations of the system, corresponding to two different theories of laws. (Both options introduce synthetic necessities, but in the next chapter I shall argue that the required concepts of synthetic necessity are unproblematic.) First, the set of predicates in the language may be augmented by adding Non-Humean ones (that is, by adding predicates that involve synthetic necessities, such as predicates that refer to bare dispositions). Call this modification of the interpretation $NH_{\Phi+}$. Second, an entirely new parameter may be introduced: a set Ω of Non-Humean synthetic necessary connections postulated to hold between predicates (members of Φ).²⁴ Call this modification of the interpretation NH_{Ω} . The basic idea is that $NH_{\Phi+}$ allows us to describe Descriptive Non-Humean possibilities whereas NH_{Ω} allows us to describe Governing Non-Humean possibilities, but more on this later.²⁵ Either way, the total number of state descriptions according to NH_{Ω} is

²⁴ I am speaking here of *predicates* rather than *properties* since state descriptions are linguistic (or propositional) entities. If Non-Humeanism is true, these connections between predicates hold in virtue of connections between properties.

²⁵ Are there reasons to prefer one Non-Humean modification to the other? I believe that neither modification suffices for the characterization of all versions of Non-Humeanism (though together they appear to do so). $NH_{\Phi+}$ appears to offer an excellent characterization of Descriptive Non-Humeanism, whereas NH_{Ω} appears to offer an excellent characterization of Governing Non-Humeanism. Here is a quick argument which suggests that NH_{Ω} entails Governing Laws: the members of Φ represent Humean properties; a necessary connection in Ω represents a necessary connection between two or more Humean properties represented by Φ ; such a necessary connection between properties constitutes a higher-order state of affairs distinct from regularities that entails regularities; hence versions of Non-Humeanism

a function of the ranges of epistemically possible values for the parameters Σ , Φ , and Ω .

Before continuing, it will be helpful to say a little more about the version of the interpretation capable of accommodating Non-Humeanism. A state description is Humean if and only if both Ω is empty and all predicates (members of Φ) are Humean. A state description is Non-Humean if and only if either Ω is non-empty or some predicate (member of Φ) is Non-Humean. The synthetic connections in Ω function in a special way. Consider the predicates ‘bachelor’ and ‘unmarried’. These predicates are logically *dependent* in such a way that every object that is a bachelor is also unmarried. The members of Ω are similar, except that the connections are not subject to analysis in the manner that ‘bachelor’ and ‘unmarried’ are; the Non-Humean connections are basic and irreducible. Thus, for example, the relation of *necessitation* (a non-symmetrical relation) is such that if P stands in the relation of necessitation to Q , any individual which instantiates P must instantiate Q on pain of inconsistency.²⁶ It seems unnecessary²⁷ to specify the nature of the function from Σ , Φ , and Ω to the number of state descriptions; what *is* necessary is to stipulate that, whatever this function is, its output is the set of *all* and *only* the maximal consistent state descriptions. Of course, this system is relativized to a language: ours will be the language generated by our experiences consistent with the methods of empirical foundationalism.

Now that our method has told us how to determine which state descriptions there are, we need to be told how to assign probabilities to them (the second issue). One proposal for assigning basic probabilities is to give equal weight to each state description:

Equiprobability of state descriptions: for any class of state descriptions that fall within a certain range of values for Σ , Φ , and Ω , the probability that a given state description is satisfied is $1/[\text{the number of state descriptions specifiable within the ranges of } \Sigma, \Phi, \text{ and } \Omega]$.

This postulate entails that every state description (that is, every epistemically possible world)

characterized by NH_Ω are versions of Governing Non-Humeanism. In the the case of Descriptive Non-Humeanism, it is the essential natures of properties themselves that (purport to) give rise to regularities. The natural suggestion, of course, is that the Non-Humean predicates added to yield $\text{NH}_{\Phi+}$ describe dispositional predicates, since predicates involving synthetic necessity in right way will satisfy the definition of dispositions introduced in the next chapter.

²⁶ Note that we get Governing Humeanism if the necessary connections in Ω turn out to be reducible or analyzable in terms of Humean-consistent entities.

²⁷ I say ‘it seems unnecessary’ because my later arguments do not require a specification of this function.

is a priori equally likely. It also entails, for instance, that each member of the set of all and only Humean state descriptions is equally likely and that each member of the set of all and only Non-Humean state descriptions is equally likely. However, Carnap (1962, 565) thought that this postulate was unsatisfactory since it appears to have the unwelcome consequence that one cannot learn from experience.²⁸ He proposed that equal weights be assigned to *structure descriptions*, where a structure description is a set of state descriptions that are structurally the same but differ only in their permutations. Thus, for example, if one flips a coin ten times, the following state description (sequence)

(7) *H H T T H T H H H T*

is one instance of the structure description

(8) *6 heads, 4 tails.*

Precisely, we have:

Equiprobability of structure descriptions: for any class of state descriptions that fall within a certain range of values for Σ , Φ , and Ω , the probability that a given *structure* description is satisfied is $1/[\text{the number of } \textit{structure} \text{ descriptions specifiable within the ranges of } \Sigma, \Phi, \text{ and } \Omega]$. Weighted probabilities are then distributed equally over the state descriptions that are members of a given structure description.

These two equiprobability postulates correspond to opposite ends of a continuum of inductive methods (the λ -continuum) that includes infinitely many intermediate postulates.²⁹ The function that specifies this continuum is monotonic, and thus if both of the above equiprobability postulates entail some proposition P , every equiprobability postulate along the continuum entails P .³⁰ Rather than arguing for the correctness of a particular equiprobability postulate, I shall

²⁸ He proved that this result held for his version of the interpretation; however, it is not at all clear that the same result holds for the version explained here since the possibility of governing laws (or, for that matter, of Non-Humean properties) may be relevant to the problem of induction.

²⁹ Explained in (Carnap 1952); see also (Howson and Urbach 1993, 66–72).

³⁰ See Michael Tooley's explanation in (Plantinga and Tooley 2008, 145–6). Note: Since these methods ultimately concern the degree of support that evidence confers upon a hypothesis, I do not think that the complication introduced by the parameter Ω affects the monotonicity of this function.

assume that the premises of a metaphysical argument must follow from both equiprobability of state descriptions and equiprobability of structure descriptions. If this can be accomplished, a fully specified interpretation of probability is unnecessary, and we are thus part way towards avoiding the charge that an equiprobability postulate can only be selected on the basis of intuition or synthetic a priori reasoning. (Note: I have come to believe that these equiprobability postulates cannot satisfactorily handle probabilistic laws; I discuss the consequences of this in 5.6. I do not discuss this worry in this chapter because the required equiprobability postulates will be very much in the spirit of Carnap's λ -continuum.)

2.4.4.3 Competing Interpretations

In this subsection, I shall very quickly discuss some of the main competitors to the logical interpretation of probability and argue that the logical interpretation provides the best method of assigning objective epistemic probabilities to metaphysical hypotheses. These arguments are not intended to be conclusive. A proper treatment of these issues would require (at least) its own book. However, I hope to suggest at the very least that my approach is promising. There are four primary interpretations of probability in addition to the logical interpretation.

Like the logical interpretation, the *classical interpretation* of probability holds that probabilities may be assigned by examining the space of possibilities through logical and semantical analysis. This interpretation is best understood as a narrow version of the logical interpretation that accepts equiprobability of state descriptions (by way of accepting a version of the *principle of indifference*), and so will not be considered further.

Frequency interpretations attempt to assign probabilities by measuring actual frequencies. For example, if one wants to know the half-life of a radioactive element, one simply observes the rate of decay over the long run. If one wants to know the probability that an average American will contract heart disease between the ages of 40 and 60, one simply looks to see what proportion of Americans (in recent times) have suffered from heart disease between the ages of 40 and 60. Unfortunately, frequency interpretations are ill-equipped to deal with metaphysical hypotheses for

the simple reason that many metaphysical hypotheses concern unobservable phenomena. Moreover, many observations relevant to the sorts of metaphysical hypotheses in question constitute *single-case uniformities*: for example, the observation that the world is full of natural regularities. But there are well-known problems for frequency interpretations associated with single-case uniformities. For example, the sequence of hands dealt in every game of poker in history, *Poker Sequence*, would appear to be a single case uniformity. Poker Sequence is being treated as a single event, and there are no other events of that type. Thus the frequency with which Poker Sequence occurs in events of the type of Poker Sequence is 1. But clearly the objective epistemic probability of Poker Sequence is much less than 1. Hence the frequency interpretation is ill-equipped for cases such as these.

Propensity interpretations hold that probability is to be explained in terms of (usually) physical properties of objects in the world. For example, the fact that the probability of a certain coin flip coming up heads is (approximately) .5 is to be explained in terms of a certain propensity or disposition of the coin. However, an account of these propensities or dispositions is required in order to make sense of this proposal. If propensities are reducible, the propensity interpretation appears to suffer from the same problems as the frequency interpretation, since a reductive account of propensities will ultimately be grounded in actually observed frequencies.³¹ If propensities are irreducible, then the interpretation has taken a serious stance on an important metaphysical issue: namely, it assumes a Non-Humean ontology. Such an assumption is inappropriate for our purposes.

Of the major interpretations, only the *subjective interpretation* remains. In my opinion, this is the most serious competitor to the logical interpretation. The general idea behind the subjective interpretation is this: there is one basic rule that one must obey when assigning probabilities, and that is to obey the theorems of the probability calculus (see (Howson and Urbach 1993) for an important defense of subjectivism). The reason that this rule has a very special status is that it is

³¹ Matters are rather more complicated than this, since there are different ways of reducing propensities; here, the relevant reduction involves a reduction to Governing Humean laws of nature. However, since I will later argue that Governing Humeanism is incoherent (in Chapter 5), I shall ignore this possibility.

required for one to have consistent beliefs. Assigning probabilities in a way that fails to respect the theorems of the probability calculus is just as bad as believing propositions that fail to respect the theorems of classical logic. But how does the subjectivist then assign probabilities? At this point, no interpretation has been offered at all.³²

In my opinion, there is no reason to accept the subjective interpretation of probability (insofar as we are concerned with metaphysical hypotheses) unless it can be shown that there are serious problems with the logical interpretation. I say this for a number of reasons. To begin, let's recall that the logical interpretation does two things: it tells us how to determine the epistemic possibilities, and it tells us how to assign probabilities to those epistemic possibilities.

First, the logical interpretation provides such a good model for describing epistemic possibilities that any competitor (in this respect) must be incorrect. Imagine an interpretation according to which the following state description wasn't considered to be an epistemic possibility: the world is exactly as it currently is, except that this chapter contained a(n additional?) grammatical error. On what grounds could such a state description be ruled out? It seems that there is none. We certainly don't want an interpretation that simply fails to recognize epistemic possibilities, and it is perfectly reasonable to think that the logical interpretation succeeds in its first task—that is, in describing the epistemic possibilities for a given language.

Second, once we have our epistemic possibilities, we want our equiprobability postulates to be as general as possible. For example, we don't want to accept an interpretation that simply ignores perfectly respectable epistemic possibilities by assigning them an epistemic probability of zero. That would amount to the claim that those epistemic possibilities aren't really epistemically possible. The logical interpretation (as I have presented it) respects this criterion by including fully general equiprobability postulates. This is desirable because, in my view, accepting a basic equiprobability postulate is very closely related to accepting a basic axiom for a system of logic. Our justification for accepting, for example, the principle of indifference (in the form of

³² It should be noted that nothing prevents the subjectivist from using the basic method of assigning probabilities specified in 2.4.4.2.

equiprobability of state descriptions) will be very similar to our justification for accepting the law of non-contradiction. If the assumption that we have justification for accepting a system of logic isn't troubling, the assumption that we have justification for accepting a system of logical probability shouldn't be troubling either.

For these reasons, I believe that we should accept the logical interpretation unless it can be shown that it is subject to serious criticisms. There are three primary objections leveled against the logical interpretation of probability.

The first objection is related to the point just made, and it is that there is no justification for choosing one equiprobability postulate over another. This objection is partially avoided by my approach, since the method I have specified requires that arguments succeed on both extremes of the λ -continuum. We are concerned only with postulates on the continuum because those postulates concern, with a great deal of generality, the degree of support that evidence has for a hypothesis, and thus it is somewhat reasonable to suspect that equiprobability postulates outside of the continuum may be ad hoc. Matters are complicated somewhat, as there have been defenses of equiprobability postulates (or continuums including postulates) that lie outside of Carnap's λ -continuum.³³ It isn't clear to me that the defenses of such methods pose a serious threat to my approach, but I haven't investigated matters sufficiently to make an informed judgment on this matter.³⁴ Though a proper discussion of these various approaches is ultimately required for the justification of the method I defend in this project, my focus here is on more general features of the method itself, so I hope that the reader will forgive me for omitting these considerations here.

The second objection is that the interpretation leads to Bertrand paradoxes, and thus that it must be false. Here is an example of one such paradox found in (van Fraassen 1989, 303).

Cube: A precision tool factory produces iron cubes with edge length ≤ 2 cm. What is the probability that a cube has edge length ≤ 1 cm, given that it was produced by the factory? Suppose we choose *edge length* as our parameter and assume a uniform distribution. Then the an-

³³ See, for example, (Hintikka 1965, 1966).

³⁴ I hope to do so in a future project.

swer will be $\frac{1}{2}$. Suppose we choose instead *area of side* as our parameter? Then the answer will be $\frac{1}{4}$. Suppose we choose instead *cube volume* as our parameter? Then the answer will be $\frac{1}{8}$. The principle of indifference leads to three incompatible answers.

In the case of Cube (and other Bertrand paradoxes) the interpretation of probability does not yield a single determinate verdict. That is what is taken to be paradoxical. But what are the implications of this paradox? Does it show that the (version of the) logical interpretation (in question) is false, or merely that it isn't well-defined for certain questions?³⁵

I think that the latter is the correct response: Bertrand paradoxes merely show that the relevant notion of logical probability is not well-defined in certain contexts. I say this because it clearly *is* well-defined in others, and has a celebrated history in the application to certain domains, such as classical games of chance. Bertrand paradoxes arise when we lack sufficient information relevant to the case at hand—when the equiprobability postulate in question can be applied in multiple ways corresponding to different variables. (The paradox wouldn't arise in for Cube if we knew the mechanism by which the machine selects cube size.) That's why, in such cases, the interpretation of probability fails to yield a single verdict. If our question is such that we possess all of the necessary information, Bertrand paradoxes just don't arise. Thus, if we can avoid such scenarios, logical probability will be well-defined.

If we are careful about the questions we ask, we can ensure that we always possess sufficient information. The central idea, which I hope to defend more carefully elsewhere, is that the interpretation of probability I have set out is properly analogous to a classical game of chance. For example, we all agree that Bertrand paradoxes do not arise for the application of the classical interpretation of probability to games of cards. Consider a deck of cards. Each card is determinately classified by two variables: suit and number. There are no other basic ways of classifying a given card; all other ways of classifying a given card are derivative on the basic criteria of classification. The number of cards is a function of suit and number. Since there are four suits and thirteen

³⁵ Some have argued that neither is the case, since such paradoxes can be avoided by more precise statements of the equiprobability postulate in question. For a recent example, see (Mikkelson 2004).

numbers, there are exactly fifty-two possible cards in a true deck. Once we've described a card in terms of both suit and number there is nothing more to say about its intrinsic nature. Further, we know exactly which cards there are.

Now consider some possible hypotheses: *the next card will be a spade*, *the next card will be an ace*, and *the next card will be the ace of spades*. These can be rephrased as follows: *the next card will be a member of the set of spades*, *the next card will be a member of the set of aces*, and *the next card will be a member of the singleton set whose member is the ace of spades*. Loosely speaking, we can identify hypotheses with the relevant sets.

Why don't Bertrand paradoxes arise in applications of the classical interpretation of probability to games of chance? There are two features:

(9) The basic means of describing each card leaves out no information (it provides a *complete* description of the card in question).

(10) We know exactly which cards there are.

Put (9) and (10) together, and there's no way to restrict information in a way that gives rise to a Bertrand paradox. Every hypothesis—every way of describing the "world" (the state of the deck)—depends on (9) and (10). The result is that there isn't the sort of missing information that gives rise to Bertrand paradoxes.

The central idea is to show that the position of the empirical foundationalist is analogous to our position with respect to games of chance. Consider the set of state descriptions for a given language L . Each state description is determinately classified by its assignments of individuals, properties, and modal connections. There are no other basic ways of classifying a given state description; all other ways of classifying a given state description are derivative on the basic criteria of classification. The number of possible state descriptions is a function of Σ , Φ , and Ω . Once we've described a state description in terms of these parameters there is nothing more to say about its intrinsic nature. Further, employing our interpretation of probability we know exactly which state descriptions there are.

Now consider some possible hypotheses: *the world is Humean* and *the world is Non-Humean*. These can be rephrased as follows: *the actual world is described by a member of the set of Humean state descriptions* and *the actual world is described by a member of the set of Non-Humean state descriptions*. Loosely speaking, we can identify metaphysical theories with the relevant sets.

My suggestion is that the set of state descriptions is perfectly analogous to a deck of cards. Above we had (9) and (10). Compare them to (11) and (12).

(11) The basic means of describing each state description leaves out no information (by definition, state descriptions are *complete*).

(12) We know exactly which state descriptions there are. (for language L)

More needs to be said about L , obviously, to show that (12) and (10) are analogous.

As it turns out, this is in principle easy to do. The language used to describe the case of the cards is one possible candidate for L . I have suggested that the language L relevant to metaphysics is a language with basic terms (those specified by our three parameters) and recursive rules. (For the moment, whether it is empirically-respectable does not matter.) It may be difficult for beings like us to determine the nature of each individual state description, and even more difficult to determine which state descriptions there are. However, it is possible in principle. Since it is possible, our position is like that in the card scenario. Since Bertrand paradoxes don't arise there, they don't arise here either. That's the idea anyway. While it warrants more careful scrutiny than I have provided here, I hope that the above is sufficient to motivate the claim that, if we are careful, we can avoid Bertrand paradoxes. As in the case of classical games of chance, the idea is that there won't be the sort of missing information that gives rise to the paradoxes.

However, sometimes it is assumed that Bertrand paradoxes arise for *every* possible question, that there are *always* different variables to which our equiprobability postulates can be applied, that logical probability is *never* well-defined. This objection is based on Goodman's new riddle of induction (see (Goodman 1955)). The crucial version of this objection is discussed in Chapter 5, so I'll wait until that point to discuss it in any detail, but for now I'll just say that it depends

on a way of defining basic predicates which allows for the following condition to obtain: the equiprobability postulates can be applied with respect to multiple (and differing) variables such that the relevant concept of logical probability is not well-defined. My response won't be surprising: Goodman's predicates are unsuitable from my methodological approach; we can reject them in a way that allows for logical probability to be well-defined for the questions we are interested in asking. But again, I'll wait until Chapter 5 to explain this response in any detail.

The third objection, which is closely related to the second, is that there is no principled way to select a language that specifies the space of epistemic possibilities. For example, Howson and Urbach (1993, p. 70) have argued that

elementary possibilities are elementary only relative to some language, and language is a human artifact whose ultimate categories stand on a footing of equality only as a result, therefore, of a collective decision that they should do so, a decision which may consequently be revoked.

The problem with this position is that we do have a principled reason for selecting a language, and thus a space of epistemic possibilities. It is specified by our empirical foundationalism, and thus is not dependent on a collective decision to accept some language over another. Within this epistemological framework, the logical interpretation properly specifies the space of epistemic possibilities. They are objective in the sense that the probabilities do not depend in any way on the particular individual that conceives of them; they are completely specified by the language. It seems that the objection would be effective if the claim were that the logical interpretation properly assigns *objective chances*, but since it is concerned with *epistemic probabilities* the objection loses force.

In sum, I do not believe that the objections to the logical interpretation are successful. I have not offered anything like a complete vindication of the logical interpretation—that would be a monumental task—but I have attempted to provide some *prima facie* justification for accepting it. For the remainder of this project, I shall assume that the logical interpretation as I have explained it (along with its claim to be offering an objective characterization of epistemic probabilities) is correct.

At this point, it will be helpful to review the progress thus far. In 2.4.2, I argued that the metaphysical hypotheses of this project require theoretical analysis, and I then explained the method of providing such analysis. In 2.4.3, I outlined a method of confirmation that, in principle, allows for the confirmation of hypotheses (even those containing theoretical vocabulary). In 2.4.4, I explained how probabilities are to be assigned to the sorts of hypotheses that are in need of confirmation. These three steps constitute the foundational methods of this project.

2.5 The Method

With the above resources in hand, I can now explain my method and show how it is immune to the sort of anti-metaphysical argument presented in 2.3.

Step 1: Apply the **first method**, the relevant principles of empirical foundationalism, to ensure that we stand in a privileged epistemic relation to all of our *basic O*-concepts (namely, those which are analytically basic).

This step provides us with the appropriate epistemological foundation. It is this step, and only this step, that implies that concept empiricism is correct. If one accepts doxastic empiricism but rejects concept empiricism, all that is required is a slight modification of this step so as to allow for basic descriptive concepts *not* given through immediate experience. Such a change has no impact on the method's ability to satisfy the requirements of doxastic empiricism.

Step 2: Apply the **second method**, the empirical foundationalist method of non-theoretical analysis, to ensure that we stand in a privileged epistemic relation to all of our *complex O*-concepts.

Together, Steps 1 and 2 are required to explain our privileged epistemic relation to our ordinary—that is, non-metaphysical—concepts; this privileged relation allows us to analyze logical and semantical connections holding between *O*-concepts.

Step 3: Apply the **third method**, the Ramsey/Carnap/Lewis method of defining theoretical terms, in the construction of all metaphysical theories to ensure that unobservable entities are

given a fully relational analysis using empirically respectable terms as specified by Step 1 and Step 2.

This step ensures that theoretical concepts are empirically respectable, and thus ensures that we have maintained our privileged epistemic relation to the content of our theoretical concepts. Just as importantly, it allows for the possibility that different theories have different observational consequences (usually given in probabilistic terms). This entails that if we have a method of genuine confirmation and suitable observations then we can be justified in believing that some theories are more likely to be true than others. The consequence of this conditional is paradigmatic of a genuine metaphysics. Thus, the method satisfies both concept empiricism and doxastic empiricism: all concepts generated by the method are empirically respectable, and it is ultimately the comparison of the observational consequences of a theory with actual observations that provides the synthetic bridge between appearance and reality.

Step 4: Define theories in such a way that the theories under consideration are mutually exclusive and jointly exhaustive.

This step places a constraint on the process of theorizing (defined in Steps 1, 2, and 3) and is required for the later application of the method of confirmation.³⁶ It is this step that ensures that the following relationship holds between the world of concepts (appearances) and the world itself (reality): it is impossible for there to be a way the world could be that does not fall under one of our theories; there are no gaps, no missing theories, and thus whatever the conclusion is, it *must* apply to the world outside of our concepts.

Step 5: Utilize the **fifth method**, the logical interpretation of probability, (i) to assign a priori probabilities to the theories in question, and (ii) to assign conditional probabilities of relevant possible observations given each theory.

For some metaphysical disputes, further steps may not be required. For example, there may be non-probabilistic cases where some theory in question is determined to be analytically false by (i)

³⁶ This step ensures that the method avoids Popper's worries about confirmation, and, in addition, the Quine-Duhem problem.

of Step 5. In probabilistic cases, both (i) and (ii) are required for the application of the method of confirmation.

Step 6: Apply the **fourth method**, the method of confirmation, to determine the probabilities of the theories conditional on each of the relevant possible observations.

The method of confirmation, of course, is content-neutral and does not assume a solution to the problem of induction. It is part of the logical system already in place. Taken together, Steps 1 through 6 establish analytic linking principles of the form $P(T|O)$, where T is the proposition that a given metaphysical theory is true and O is a(n) observation statement(s).³⁷ The final step is as follows:

Step 7: Observe! Then compare the observations with the linking principles to determine the a posteriori probabilities of the theories under consideration.

Together with the preceding steps, the result is as follows. The empirical foundationalist can practice metaphysics so long as the following conditions are met: a proper defense of the relevant interpretation of probability can be provided, and Steps 1 through 7 are satisfied. If empirical foundationalism does not require synthetic a priori reasoning then neither does this method as a whole.

Each step is required for the method as practiced by the empirical foundationalist.³⁸ Without Steps 1 and 2, the method would fail to be empirically respectable. Without Step 3, metaphysics would merely be conceptual analysis. Without Step 4, the method of confirmation would not be possible, and there would be no guarantee that the method tells us something about the world itself (that is, about reality rather than mere appearance). Without Steps 5 and 6, the scope

³⁷ Readers less familiar with conditional probability can think of the linking principles as having the form of conditionals: $O \supset W$ (where O is some possible empirical observation statement and W is some proposition about the probability that a theory about the world itself is true).

³⁸ This method is similar in some respects to Michael Tooley's basic approach to practicing metaphysics, most clearly outlined in (Tooley 1987). He endorses something very much like empirical foundationalism, he appeals to the Ramsey/Carnap/Lewis method of defining theoretical terms, and he appeals to Carnap's logical interpretation of probability. The version defended here is intended to be an improvement over Tooley's in multiple primary respects: first, it incorporates additional elements so as to be more precise (Step 4 is the most important); second, it is constructed to allow for an explicit demonstration that the results of the method are characteristic of *genuine* metaphysics.

of genuine metaphysics would be extremely limited; one could reject theories that are internally inconsistent, or those that entail a contradiction when conjoined with some empirical observation, but that is all. Without Step 7, (at best) one could only determine the a priori probabilities of theories.

Before continuing, it may be helpful to explain explicitly how the method results in justified beliefs. We have justification for all beliefs generated in accordance with Step 1. These are just those justified non-inferentially through immediate experience. To simply, following Carnap, let's call such beliefs *protocol sentences*. We have justification for all beliefs generated in accordance with Step 2. These are just those justified inferentially through the conjunction of the system of non-theoretical analysis plus the protocol sentences. We have justification for all beliefs generated in accordance with Step 3. These are just those justified inferentially through the conjunction of the system of theoretical analysis plus those complex beliefs from Step 2 and the protocol sentences. Step 4 merely places a constraint on the way in which theories are generated using Steps 1 through 3. It yields a set of propositions B that is a subset of those that can be generated by 1 through 3. Since beliefs in the latter can be justified, so can belief in those propositions B . Step 5 merely applies a set of logical principles to B in order to assign probabilities, forming a set of beliefs about probabilities P . Since these principles are justified, and beliefs about B are justified, the resulting beliefs about P are justified. Step 6 merely applies the theorems of the probability calculus to P , transforming the relevant beliefs into our analytic linking principles. Since we are justified in accepting these theorems, and justified in accepting P , so we are justified in accepting the linking principles. Step 7 merely takes the protocol sentences from Step 1 and applies them to the linking principles of Step 6 in order to derive rational beliefs about the world given one's total evidence. If we thought of the linking principles as conditionals rather than as atomic statements of conditional probability, this would be a straightforward application of modus ponens. Since the linking principles are actually atomic statements of conditional probability, matters are somewhat more complicated. The linking principles state the probability of metaphysical theories given *possible* observations; that is, a given principle effectively says "if we were to observe O , the posterior

probability of theory T given O would be k ." The final step tells us what our actual observations are; it tells us which linking principles are relevant; it tells us which principle to accept as providing the posterior probability of the theory given *actual* observations. Though technically different from an application of modus ponens, the effect is the same as far as justification is concerned.

This is a very interesting result. Contrary to Hume, empirical foundationalism (British empiricism, more or less) does not lead to metaphysical skepticism. In principle, immediate experiences *are* relevant to metaphysical theories. If we reject empirical foundationalism because we reject empiricism, a subtle modification of the method—one which makes the relevant changes to Steps 1 and 2, plus adds any conditions necessary to weigh the conclusions of this method against those provided by any other relevant method of practicing metaphysics—will yield the result that metaphysics can be informed by immediate experience. Again, a very interesting result.

2.6 Conclusion

I have argued that, contrary to conventional wisdom, empirical foundationalism (which is essentially British empiricism) does not lead to skepticism of matters metaphysical. Further, I have shown precisely how empirical foundationalists can practice metaphysics in a way that assures that their conclusions will be about the world itself (reality) rather than our concepts of the world (appearances). In Chapters 3 through 6 I shall argue that our observations of the physical world constitute strong evidence in favor of Governing Non-Humeanism. These chapters may be considered a sample application of the method, but their arguments are sufficiently general that they may be adopted by philosophers rejecting many of the central epistemological and methodological assumptions of this project. It is my hope that, in part, the success of those arguments will convince philosophers that the method defended in this chapter is promising.

That being said, here is where I think the method currently stands. Though it is promising, it is not yet fully vindicated. It must be demonstrated that we are justified in believing some particular system of logic to be correct, and I do not know whether this can be achieved in a manner consistent with the empiricist spirit of empirical foundationalism. However, as mentioned earlier,

I am willing in principle to allow for the possibility of synthetic a priori reasoning as applied to issues such as these; I think, for example, that the law of non-contradiction is rather more privileged than, for example, the principle that all causes are prior to their effects. Further, more work is required to demonstrate that the logical interpretation of probability is correct (or, at least, that there is some correct method of assigning objective epistemic probabilities). And, of course, empirical foundationalism must be established as the correct epistemic theory (no small task here, either). I have not satisfied these burdens here (though I hope to do so elsewhere). If these burdens can be satisfied, metaphysics is possible without synthetic a priori reasoning. If they cannot, the method is still of great value. It clearly demonstrates that metaphysics can be empirically informed by immediate experiences. Even if we do possess synthetic a priori reasoning, metaphysics need not be an exercise in synthetic a priori reasoning *only*.

2.7 Appendix: Miscellaneous Methodological Criteria

The purpose of this section is to discuss two common methodological criteria and explain their relevance to my methodology.

2.7.1 Ontological Economy and Occam's Razor

Many philosophers have thought that *ontological economy*—the degree of ontological simplicity of a theory in terms of both types of entities (*quality*) and number of entities (*quantity*)—is closely related to the likelihood that a theory is true: simple theories (those with fewer entities or types of entities) are thought to be more likely to be true than more complex theories. Or, to put the matter simply, some philosophers have thought that theories involving robust ontologies are simply unlikely to be true. This, I believe, is a mistake. Though there are clear-cut cases in which changes in ontological economy (and nothing else) affect the a priori probability of a theory, it is a mistake to generalize on the basis of these cases that ontological economy itself ought to be privileged. I shall argue that we needn't try to formulate some principle concerning ontological economy; the method above already takes this matter into account.

I shall start by offering an example.³⁹ Let R be the proposition that there are cosmic regularities. Let G be the proposition that there are governing laws understood as states of affairs that entail R . The method entails that the a priori probability of $R \& G$ is less than that of R , since R could be satisfied even if G were not the case. (After all, R is equivalent to $(R \& G) \vee (R \& \neg G)$.) The relevant difference between G and R is that G contains a more robust ontology. This sort of example has led some philosophers to hold that robust ontologies are less likely to be true than sparse ontologies. The problem with this approach is that hypotheses R and G are hypotheses of different levels of generality. The hypotheses that are really in need of comparison are the following: $R \& G$ and $R \& \sim G$. These are the two hypotheses that make substantive metaphysical claims. R , on its own, does not; it is an empirical statement. G and $\sim G$ do not share this feature. Thus, a generalization from the fact that $P(G \& R) < P(R)$ to the conclusion that ontological economy affects the probability of a hypothesis simply misses the point. Through careful exploration, one might discover that a certain hypothesis allows for fewer state descriptions than its competitor. This would give one a reason to think that the former is less likely than the latter (this sort of approach will be utilized in Chapters 3, 5, and 6). But, in fact, this is not what proponents of ontological economy usually do.

(The example above concerns the economy of *quality*. The matter may be somewhat different for economy of *quantity*. However, I shall ignore this issue since most metaphysical issues primarily differ with respect to quality; the theories considered here do not differ considerably with respect to the quantity of entities they postulate.)

A justification of the preference for ontological economy requires something else entirely. But what? No doubt there are epistemological concerns. For example, Schaffer (2008, p. 96) summarizes (though he ultimately rejects the argument):

The idea is that our nomic knowledge is ultimately based on our observation of regularities in history, so that if laws were more than such regularities, we could have no access to this further feature. So the ontologist should just drop the further feature and limit herself to the regularities.

³⁹ I thank Michael Tooley for suggesting this particular example.

But such concerns are based on the failure to recognize the methods of determining the observable implications of metaphysical hypotheses and confirmation of such hypotheses; they simply fail to recognize the methods set out in this chapter. Therefore, this sort of epistemological worry does not justify the bias towards ontological economy. And, of course, one who endorses something like Humeanism on the basis of this argument is *not* limiting himself to the regularities; he is also endorsing a controversial metaphysical hypothesis.

The remaining justification (of which I am aware) of ontological economy is based on the idea that certain metaphysical entities are strange or mysterious (see for example (Mackie 1977) and (Sider 2001)).⁴⁰ I am sympathetic to the idea that we shouldn't utilize concepts that we don't understand. If *mysteriousness* is just a property of poorly-defined concepts, then I have no objection to this idea. But our method of defining theoretical terms and referring to theoretical entities eliminates mystery since it ensures a proper understanding of the relevant concepts. If mysteriousness is something else, then it is very difficult to see what can be said in defense of ontological economy.

To sum up, if there is a role for ontological economy (and I'm not sure that there is one), it is this: ontological economy can act as a rule of thumb. If we look at two theories, and cannot say whether one has more explanatory power than the other, and don't know enough about them to try to determine their prior probabilities via the method outlined in 2.4.4, then we can provisionally select the simpler of the two theories. Or, if we have a set of equally explanatory theories, one more robust than the other, we can select the sparsest one. Ultimately, however, we shall see that such situations rarely obtain. Once serious philosophical investigation takes place, we have little or no use for such rules of thumb. The method of assigning probabilities already accomplishes exactly what we should like to accomplish. There is no need—and no room—for unmotivated additions that conflict with the method.

⁴⁰ Mackie's prejudice is against *normative* properties; Sider's is against *modal* properties.

2.7.2 The Status of Modal Intuitions

Some philosophers believe that we have special epistemic access to matters metaphysical—that we have certain *intuitions* (or *seemings*) that are reliable indicators of the way the world is. The relevant intuitions here are not intuitions that a certain object or class of object “falls under” a certain concept; rather, they are intuitions that a given theory or proposition is *true* or *possible* or *necessary*, et cetera. Modal intuitions are of particular relevance to this project since most explanations of natural regularities include modal elements. For example, we might ask ourselves whether it is possible for there to be uninstantiated laws of nature—that is, a law for which there is no corresponding regularity. Some philosophers (Tooley (1977; 1987) and Carroll (1994)) believe that it is possible.⁴¹ They do so by describing state descriptions in which there is such a law that seem (to them at least) to be possible. Others (Beebe (2000)) believe that uninstantiated laws are not possible. Whatever the arguments of these philosophers, our question is just whether the presence of a modal intuition that *P* provides one with epistemic reason to believe that *P* is the case.

The appeal to modal intuitions is sometimes a veiled way to invoke synthetic a priori reasoning. Many of our intuitions concern synthetic principles: for example, the principle that a given state description *S* is possible. But what is the sense of possible being used here? We generally have good epistemic access to the matter of whether something is epistemically possible. But do we have good epistemic access to the matter of whether something is metaphysically possible—that is, *genuinely* possible, possible *simpliciter*, possible *in the world itself*? To make the latter claim is to say that one has some reason for believing a synthetic principle with no empirical evidence. In other words, it is to accept synthetic a priori reasoning. It would be nice if our intuitions could be granted some special epistemic status, but I’m not sure how such status can be granted. Of course, it may be that they provide the only possible justification for e.g. believing that one system of logic is more likely to be correct than another. However, I believe that our modal in-

⁴¹ I have in mind Tooley’s ten-particle case and Carroll’s mirror case. To be fair, it isn’t clear that Tooley intends for the operative sense of possibility to be *metaphysical* possibility, but this seems to be a way in which many philosophers view cases like these.

tutions concerning cases such as these are rather different from our modal intuitions concerning the metaphysical possibility/impossibility/necessity of complex state descriptions—the former are less susceptible to problems of intersubjectivity than the latter. Thus, in this project I shall do my best to avoid appeals to intuition entirely, while acknowledging that my project may require them at the foundational level.

Chapter 3

Non-Humeanism

3.1 Introduction

In the empiricist tradition, necessity is often given the following reductive account: a proposition is *necessarily true (false)* if and only if it is true (false) in virtue of its meaning or logical form—that is, if and only if it is *analytically* true (false). This definition is consistent with empirical foundationalism in that the resulting form of necessity is perfectly intelligible even to the strict empiricist. However, this account entails that there is no synthetic necessity in the world; it is the paradigm statement of Humeanism. On this reductive account, propositions such as $2+2=4$, $p \vee \sim p$, and *no bachelors are married* are necessarily true, but propositions such as *all objects with mass are attracted to one another* and *light travels faster than anything else* are not. The propositions that are necessarily true according to Humeanism are interesting, but they don't seem to be the sorts of necessary truths that are relevant to natural regularities; that is, they don't seem to have explanatory power with respect to the regularities of the natural world. Thus, the reductive account of necessity seems to rule out any explanation of regularity that entails the following schema: regularities occur because they *must*. (Whether or not it does in fact rule out such explanations is, for the moment, an open question; it will be examined carefully in Chapter 4, the chapter on Governing Humeanism.) This sort of explanation seems to require a different sort of necessity altogether: synthetic necessity—irreducible necessity that resides in the world itself. We must determine whether this sort of necessity is intelligible.

As discussed in the last chapter, irreducibly modal features of the world are not immediately

observable. Hence, if the concept of synthetic necessity is to be analyzed at all, it must be given a theoretical analysis. The first goal of this chapter is to demonstrate that the concept of synthetic necessity (and thus Non-Humeanism) is intelligible according to empirical foundationalism. The second is to argue that the postulate of synthetic necessity has no negative effect on the a priori probability of a theory. In short, this chapter offers a preliminary defense of Non-Humeanism, specifically Governing Non-Humeanism.

This chapter is organized as follows. 3.2 constitutes my analysis of synthetic necessity. I start this section by offering an analysis of the concept of *metaphysical possibility*. Humeanism and Non-Humeanism are then distinguished with respect to their treatment of metaphysical possibility. 3.3 considers a priori objections to a well-known version of Non-Humeanism, the view that there are bare dispositions. I find these a priori arguments inconclusive. 3.4 considers the a priori probabilities of Humeanism and Governing Non-Humeanism, and it is argued that Governing Non-Humeanism is a priori just as likely as Humeanism.

3.2 A Theoretical Analysis of Synthetic Necessity

The following subsections provide the theoretical analysis of synthetic necessity. The first provides the stipulative analyses of metaphysical possibility and synthetic necessity. The remaining subsections consider implications of the analyses and respond to issues brought up by the analyses.

3.2.1 The Analysis of Metaphysical Possibility

Some philosophers (van Inwagen 1998) define *metaphysical possibility* as possibility *simpliciter*, or possibility *period*. In Chapter 2, I explained how the logical interpretation of probability provides a concept of epistemic possibility, which may be more carefully developed as follows. Logic gives us a notion of *formal consistency*. A formal system implies that contradictory sets of propositions are impossible, but, since it lacks content, it doesn't do much else for the pursuit of metaphysics. This problem is remedied by adding meanings to terms in a formal

language. This is the role of a semantics, which gives rise to a new concept: *semantical consistency*. For example, in natural English the proposition *the cat is on the mat* is semantically consistent, but the proposition *the bachelor is married* is semantically inconsistent. *Logical/semantical consistency* is just the intersection (the overlap) of formal and semantical consistency; a proposition is logically/semantically consistent if and only if it is both logically consistent and semantically consistent. Different languages will obviously yield different sets of logically/semantically consistent propositions, but the assumption here is that empirical foundationalists will apply their own language as generated by methods 1 through 3 explained in Chapter 2. From this discussion it should be clear that Carnapian state descriptions just are logical/semantical possibilities. I introduce this concept because it will be helpful to compare the analysis of metaphysical possibility with the analysis of logical/semantical consistency.

It is an open question whether metaphysical possibility (whatever that is) is the same thing as logical/semantical consistency. It is not initially obvious that the concept of metaphysical possibility is analyzable (notice I do not say reducible) within empirical foundationalism. The same goes for the concept of synthetic necessity. These concepts probably weren't analyzable within empirical foundationalism during Hume's time, since he lacked the technical resources that are constituents of the method explained in Chapter 2. My goal is to show that these technical resources do in fact make the crucial difference: they allow for a theoretical analysis of metaphysical possibility and synthetic necessity that can be offered in empirically respectable terms.

I wish to emphasize three important points before continuing. First, it should be understood that my analyses of metaphysical possibility and synthetic necessity are *stipulative*; I do not claim that this is how other philosophers have understood these concepts (though I do think that it is how many of them have understood them). Second, the goal of section 3.2 is merely to show that the claims of Non-Humeanism are perfectly intelligible to the empiricist; the goal is *not* to argue that metaphysical possibilities are distinct from logical/semantical possibilities; the goal is *not* to argue that the Non-Humean conception of metaphysical possibility is better than, or more accurate than, the Humean conception of metaphysical possibility (in fact, the concepts are the

same); the goal is *not* to argue that there *are* synthetic necessities. I am merely arguing (in this section) that Non-Humeanism is intelligible. Third, I do not here attempt to explain either the ground of synthetic necessities or our epistemic access to these necessities. The former appears to be impossible if doxastic empiricism is true; the latter is explained by the arguments of later chapters, which rely on the method set out in Chapter 2.

My explanation of the concept of metaphysical possibility is given in three steps. The first step is to provide a semantics for modal logic. The well-known possible worlds semantics will suit our purposes just fine.¹ (I assume a basic familiarity with this semantics, and thus my explanation of its basic features is extremely brief.) A statement such as ‘it is necessary that’ is analyzed in terms of *possible worlds*. What is a possible world? For the purposes of a logical system, it doesn’t matter. Philosophically, however, this is a very controversial matter. Fortunately, we already have a suitable, empirically respectable candidate for possible worlds: Carnapian state descriptions. Thus, for our purposes, a possible world is nothing more and nothing less than a logically/semantically consistent state description. The semantics uses possible worlds to define the standard modal operators, \Box and \Diamond , by appealing to an *interpretation*, where an interpretation is a specification of a set of possible worlds (for the moment, I ignore accessibility relations for simplicity). A proposition that is prefixed with \Box (is necessarily true) is true if and only if it is true at every world in the interpretation; a proposition that is prefixed with \Diamond (is possibly true) is true if and only if it is true at at least one possible world in the interpretation. The empirical foundationalist has no trouble understanding \Box and \Diamond for the simple reason that the empirical foundationalist has no trouble understanding (a) the concept of a possible world, (b) the concept of an interpretation, and (c) the terms used in the definitions of \Box and \Diamond .

For the purposes of illustration, I shall now demonstrate how an empirical foundationalist understands *analytic* necessities within the context of the possible worlds semantics. Since an interpretation is constructed out of state descriptions, and since all state descriptions are by

¹ This version of the semantics is due to Kripke (1963). For a clear and accessible introduction to possible worlds semantics, see (Forbes 1985).

definition logically/semantically consistent, no interpretation includes any impossible state descriptions. By definition, analytic necessities are true according to every state description. Hence, every analytic necessity is true in every possible state description. Hence, every analytic necessity is true in every state description in every interpretation. Hence, every interpretation makes all analytic necessities necessarily true.

Synthetic necessities, if there are any, are not like that. All synthetic necessities are, in some sense, conceivably false; they are not logically/ semantically necessary, and thus the sense in which they are necessary can only be the metaphysical sense. Hence, for every synthetic necessity, there is a logically/semantically consistent state description that entails (more precisely, *would entail* if it were part of the relevant interpretation) that it is false. If such a state description is included in the interpretation, the synthetic necessity in question will fail to hold. Thus, synthetic necessities hold only on *restricted* interpretations: those interpretations in which the epistemically possible state descriptions which constitute counterexamples to the synthetic necessity (or necessities) in question are omitted.

In light of this, the second step is to demonstrate that the empirical foundationalist is capable of understanding the concept of a *restricted interpretation*. This is very simple. A restricted interpretation is nothing more than a set of logically/semantically consistent state descriptions, so of course there is no trouble in understanding this concept. Here is a simple illustration. Suppose that one has a very simple language: this simple language permits only four state descriptions: α , β , γ , and δ . These are the logical/semantical possibilities. Now consider interpretations A and B:

A: $\{\alpha, \beta, \gamma, \delta\}$

B: $\{\alpha, \beta\}$

Both interpretations are perfectly intelligible. One can simply apply the definitions of \square and \diamond to determine which propositions are possible and which are necessary on the two interpretations. Given the nature of α , β , γ , and δ , it may turn out that B includes some necessary truths not included in A. For instance, α , β , γ , and δ may be state descriptions generated by a language

with two possible sentences, P and Q , such that $\alpha = \{P, Q\}$, $\beta = \{P, \sim Q\}$, $\gamma = \{\sim P, Q\}$, and $\delta = \{\sim P, \sim Q\}$. In this case, $\Box P$ is true on interpretation B, the restricted interpretation, but not on A, the unrestricted interpretation.

The third step is to demonstrate that the empirical foundationalist is capable of understanding the concept of an interpretation's being the actual interpretation, that is, the interpretation that "correctly describes" the world. Note that Humeanism already includes the concept of an interpretation's being actual; it holds that the *unrestricted* interpretation is actual.² Not explicitly, perhaps, but this is entailed by the Humean's claim that propositions are necessarily true *only if* they are analytic. Humeanism is, after all, a theory about the world. And so the Humean can hardly criticize the Non-Humean for including this same concept.

(There are additional reasons for thinking that the concept of a restricted interpretation's being actual is a familiar concept to us, but these are discussed in the next subsection. There, I shall also discuss the important charge that there is something wrong with the application of "actuality" to interpretations—that is, with the concept of an interpretation's being the actual interpretation.)

Here, then, is the analysis of metaphysical possibility:

MP: P is metaphysically possible =_{df} *P is part of a state description that is a member of the interpretation I that is actual.*

This definition is rather "thin." It only depends upon the concepts of *state description*, *possible worlds semantics*, *interpretation*, and *actuality*. But this is a desirable result from the perspective of empirical foundationalism. The above definition doesn't explain how metaphysical possibility is related to logical/semantical consistency, nor should it; that would be to stipulate a solution to an important philosophical issue. It doesn't explain what our epistemic access to the metaphysically possible is. It doesn't assume either Humeanism or Non-Humeanism. It leaves these matters open, which is exactly what we want, given that it is, at the moment, an open question

² In the next section, we shall see that in fact the Humean's concept of an interpretation is also restricted, but in a slightly different sense.

whether Humeanism or Non-Humeanism is correct. Additionally, it is consistent with the view that metaphysical possibility is possibility *simpliciter*.

We are now in the position to complete our analysis of *synthetic necessity*. We saw above that there are synthetic necessary truths if and only if the actual interpretation is *restricted*. Thus the empirical foundationalist can distinguish Humeanism and Non-Humeanism as follows:

Non-Humeanism is true (that is, there is synthetic necessity) =_{df} the actual interpretation is restricted.³

Humeanism is true (that is, there is no synthetic necessity) =_{df} the actual interpretation is unrestricted.

I wish to stress once more that the definition of metaphysical possibility is neutral between Humeanism and Non-Humeanism. It is important to recognize the limited nature of my claim in this section. In saying that Non-Humeanism is intelligible, I am merely saying that we all understand what the Non-Humean means when he says that a state description cannot obtain (or that it must obtain). There is nothing mysterious or confusing about the Non-Humean's claim. Though there might be something mysterious about the source and ground of the necessity posited by the Non-Humean, that is an entirely separate issue, and does not affect in any way the content of the claim itself.⁴

³ Note that Governing Non-Humeanism (NH_Ω) and Descriptive Non-Humeanism (NH_{Φ+}) provide different methods of restricting interpretations.

⁴ Robert Hanna has helpfully suggested that further work is required to show that the analysis provided here constitutes an analysis of synthetic necessity. The basic charge is that a more careful discussion of the relevant semantics is required. After all, I have said that a proposition is analytic if and only if it is true in virtue of *meaning* or logical form, but I have said little about what it is for a proposition to be true in virtue of meaning. The idea is then that different languages can include different meaning postulates, with the result that different languages can include different analytic propositions. Though I have yet to investigate this matter carefully, my hope is that the worry can be avoided by carefully specifying the conditions under which an empirical foundationalist is justified in accepting a given meaning postulate. Some meaning postulates may turn out to be analytic in virtue of being rules that merely state purely logical connections between simple concepts and more complex ones. Others may turn out to be themselves synthetic. The idea will be to distinguish these different sorts of meaning postulates, and then argue that "analytic" propositions generated by synthetic meaning postulates count as synthetic propositions in the sense we care about here. For example, suppose we introduce the meaning postulate *no shade in the spectrum of colors is a distinct shade in the spectrum of colors*. According to this postulate, the following proposition is analytic: *no shade of red is a shade of green*. However, since the meaning postulate itself appears to be synthetic, the sense in which the latter proposition is analytic isn't the same as the sense in which propositions such as $p \vee \sim p$ or *all bachelors are unmarried* is analytic.

3.2.2 More on the Concept of Actual Restricted Interpretations

The purpose of this section is to elaborate on my claim that the concept of a restricted interpretation's being actual is intelligible. To do so, I wish to point out two areas in which the concept is already familiar.

First, possible worlds semantics are usually more complicated than described above. In addition to the above components, most systems include a set R of *accessibility relations*, relations that allow for a system to capture the idea that that which is possible at one world may not be possible at another. As before, the semantics uses possible worlds to define the standard modal operators, \Box and \Diamond , by appealing to an *interpretation*. However, in this case an interpretation is a specification of a set of possible worlds *plus* accessibility relations between worlds. A proposition that is prefixed with \Box (is necessarily true) at a world w is true if and only if it is true at every world in the interpretation that is accessible from w ; a proposition that is prefixed with \Diamond (is possibly true) at w is true if and only if it is true at at least one possible world that is accessible from w .

How is this relevant? The view that the actual interpretation is unrestricted is consistent with only one characterization of R : that every world is accessible from every world. But certainly we understand systems in which R has a different set of relations. We understand exactly how they work, and what their proponents claim. And we have examples in which different characterizations of R correctly and completely describe some portion of reality. For example, the way in which it is possible for the king to move in a game of chess depends on the state of the game—specifically, the king can be castled only if there are no pieces between the king and the rook in question, and neither the king nor the rook have previously moved. This rule is perfectly intelligible. So certainly the concept of a restricted interpretation's being the actual interpretation is intelligible.

Second, those philosophers (Humeans included) who accept something like classical logic (C) (or, for that matter, pretty much any widely studied system of logic) already accept a restricted interpretation. Consider a system of classical logic L minus one fundamental rule of inference. C

includes some theorems that L does not. C allows for the derivation of contradictions not derivable in L . Thus, the set of logically/semantically consistent state descriptions in C is a proper subset of the set of logically/semantically consistent state descriptions in L . But many philosophers (empiricists included!) accept the “move” from L to C , so they must be familiar in some sense with the concept of the actual interpretation’s being restricted. *Their* concept of the actual interpretation is a restricted interpretation. I cannot see how the move from Humeanism to Non-Humeanism is any more problematic (insofar as we are concerned with intelligibility) than the move from L to C . This problem generalizes. There is nothing unique (relevant to our purposes) from the move from L to C . The same can be said for the acceptance of virtually any system of logic at all. Those philosophers who think that some logical axioms are true (regardless of what those axioms are) will be in the same boat.

Of course, we might question one’s motives in moving from L to C . There are a number of difficult epistemological worries that arise in making such a move. But none of these issues are relevant to whether or not the move is *intelligible*, and the move from Humeanism to Non-Humeanism is not problematic in any additional way. The problem of understanding the concept of an interpretation’s being actual is exactly the same for any restricted interpretation, whether the restriction is from $S5$ to $S4$, L to C , or from Humeanism to Non-Humeanism.

This argument may be summed up by the following dilemma: either we have reason to accept that a system of logic (or a set of logical principles) is correct or we do not; if we do, then we are perfectly comfortable with the concept of a restricted interpretation; if we don’t, then we cannot form any complex concepts nor can we ever have justification for any belief involving complex concepts. The reason is that, as empirical foundationalists, in order to form complex concepts and justify inferentially justified beliefs we need to be justified in believing that our system of analysis is truth-preserving. Our system of analysis is our system of logic. Thus rejecting the concept of a restricted interpretation’s being actual has disastrous skeptical consequences. The only recourse for those who reject it is to adopt a thoroughgoing pragmatism. This is an option that must be dealt with at some point, but for this project I have assumed that empiricists can be

justified in their logical beliefs. Here I have demonstrated that if empiricists can be justified in their logical beliefs, they are also capable of possessing the concept of synthetic necessity.

3.2.3 Synthetic A Priori Reasoning and Restricted Interpretations

Suppose for the moment that we possess a faculty of rational intuition that allows for synthetic a priori reasoning. Through the use of this faculty, we come to believe that *necessarily, no shade of red is a shade of green*. Of course, this proposition entails that the actual interpretation is restricted since (according to conventional wisdom) the proposition is not analytic. The actual interpretation doesn't include any state descriptions in which a shade of red is a shade of green, despite the fact that such state descriptions are logically/semantically consistent.

If we are able to entertain this thought experiment, what reason can we have for thinking that the concept of a restricted interpretation's being actual is unintelligible? My point is merely that the claim that Non-Humeanism is unintelligible has far-reaching consequences.⁵ Here, it seems to suggest that we do not even understand the debate between *rationalism* and *empiricism*.

3.2.4 An A Priori Argument Against Humeanism?

There is an a priori argument that attempts to show that Humeanism is contradictory. My reasons for including it in the section concerned with the intelligibility of Non-Humeanism will become apparent shortly. Consider again the Humean's account of necessity: a proposition is *necessarily true (false)* if and only if it is true (false) in virtue of its meaning or logical form. Our question is whether this account of necessity is *necessarily* true.

If Humeanism is the position that the above account of necessity is necessarily true, then Humeanism is contradictory. Why? The proposition *only analytic propositions are necessarily true* is synthetic, and thus self-contradictory.

If Humeanism is only contingently true, then Humeanism allows for the possibility of synthetic necessities. But the position that there are no synthetic necessities, though there could be, is

⁵ Readers who find this sort of consequence unappealing will also be interested in Carroll's (2004, Chapter 1) centrality thesis, according to which almost all of our everyday concepts have nomological commitments.

a bit strange—perhaps unintelligible. There is, however, a way to make sense of it. But it involves accepting a restricted interpretation—one in which there are worlds not accessible from the actual world but accessible from some possible worlds. For example, suppose that w is the actual world, and that w , u , and v are worlds accessible from w , and that only w and u are accessible from u , and that the result is that some synthetic proposition P is necessarily true in u (for example, because P is true in both w and u but not v). v is not accessible from u , but it is accessible from w , and hence P is not necessarily true in w though it is necessarily true in u .

The upshot? Humeanism, if it is to be intelligible, includes the concept of the actual interpretation's being restricted.

In this section, I have merely tried to show that we understand the claims of the Non-Humean, and it is important to recognize that understanding the claims of the Non-Humean does not entail understanding the “source and ground” of synthetic necessity. Whether we have any reason to accept Non-Humeanism is the project of the final section of this chapter and the later chapters of this project.

3.2.5 Implications and Further Discussion

I have not actually given a complete, positive empirical analysis of the relevant concept of actuality. However, I have shown that it is a concept with which most empiricists are familiar. Consider some arbitrarily chosen state description S . A Non-Humean might explain his position as follows:

S is perfectly intelligible, perfectly consistent. It is certainly conceivable. Nonetheless, there is a theory T that says that S just cannot obtain. I don't claim to know what it is about S that makes it the case that it cannot obtain. I'm not trying to explain what sort of reasons we might have for accepting T . I'm just saying that we understand T . That's all.

I don't think that there's anything confusing about this position for the same reason that I don't find anything confusing about the following claim: there is a system of logic with certain axioms, and it turns out that some of those axioms are false. Of course, we might wonder what reasons

one could possibly have for accepting it, but we don't, I claim, find it unintelligible. The analysis implicit in this quote is a *negative* one. It merely suggests a necessary condition for something's being metaphysically possible. But we understand it. And that is enough for our purposes.

The analysis of synthetic necessity provided above does have some interesting implications. Since synthetic necessity is not mysterious, we can entertain conflicting synthetic necessities. For example, we can entertain both

(1) Necessarily, *nothing* travels faster than the speed of light.

and

(2) Necessarily, *something* travels faster than the speed of light.

Of course, (1) and (2) are logically inconsistent. Though they are both epistemically possible (we can understand entertain both propositions), they cannot both be metaphysically possible. The upshot is that we can specify different versions of Non-Humeanism, many of which are inconsistent with one another. However, as long as each version offers an internally consistent restriction of the interpretation, each is logically/semantically consistent.

This phenomenon is by design. We want our analysis of synthetic necessity to ensure that the set of logical/semantical possibilities is a superset of the set of metaphysical possibilities in order to justify our inferences from the former to the latter. The above account helps to ensure that this condition is met.

Another interesting implication of this account is that it allows for the possibility of distinct state descriptions that are categorically equivalent. Roughly, the idea is that two state descriptions are categorically equivalent if and only if they have all the same non-modal features, but this term is defined most perspicuously through reference to features of the logical interpretation of probability set out in Chapter 2. Two state descriptions are *categorically equivalent* if and only if their assignments of individuals (from parameter Σ) and Humean predicates (from parameter Φ) are the same and they differ only with respect to their Non-Humean predicates (from Φ) or

their Non-Humean connections (from parameter Ω). Thus, for example, the following two Non-Humean state descriptions are categorically equivalent, and each is internally consistent.

(3) Fa, Ga, Fb, Gb , and F -ness *necessitates* G -ness.

(4) Fa, Ga, Fb, Gb , and G -ness *necessitates* F -ness.

These examples are versions of Governing Humeanism since they are characterized in terms of NH_{Ω} , but similar examples can be given using the other characterization. One virtue of this characterization is that it offers a very clear picture of the fact that state descriptions can be categorically equivalent but modally distinct.

I have argued that empiricists can provide an analysis of Non-Humeanism. The question now is whether there is any reason to accept a Non-Humean ontology over a Humean ontology. In 3.3, I examine three arguments against a prominent version of Non-Humeanism: the view that there are bare dispositions. Even if Governing Non-Humeanism can dispense with bare dispositions altogether (whether this is so is not clear to me), the arguments against bare dispositions are worth considering since they rely on the modal features of bare dispositionalism. I find that these arguments are inconclusive at best. In 3.4 I consider the a priori probabilities of Descriptive Humeanism and Governing Non-Humeanism, and I argue that we have reason to think that, prior to any observation at all, Governing Non-Humeanism is just as likely to be true as Descriptive Humeanism.

3.3 Bare Dispositionalism: A Type of Non-Humeanism

Consider the following very basic explanation of a natural regularity. All salt is such that, if it were placed in water, it would dissolve. Why is this? Salt has the property *water-solubility*, which is to say that there is a Non-Humean connection between salt, water, and solubility. That connection is the ground of the subjunctive conditional in question.⁶ Now, subjunctive conditionals are often used to distinguish between *dispositional* and *categorical* properties. The traditional way of

⁶ I shall consider this sort of explanation with much greater care in Chapter 6.

drawing the distinction is as follows: dispositional properties entail subjunctive conditionals; categorical properties do not. (Note that I'm going to offer a different distinction.) Since subjunctive conditionals have modal character, the view that there are *irreducible* dispositions constitutes a version of Non-Humeanism. The purpose of this section is to consider a priori arguments against this version of Non-Humeanism (since they are common and well-known) and argue that they do not give us reason to reject Non-Humeanism.

Before beginning, I shall make a few notes on my requirements for a theory of dispositional properties. Since the purpose of this project is to provide an explanation of natural regularities, I am (here) only interested in dispositional properties insofar as they are relevant to this project. I shall ignore many interesting and otherwise important distinctions, such as the distinction between dispositional and categorical *predicates*. In short, I shall define dispositions in a way that suits my purposes. This might sound dogmatic. However, here one must be mindful of my guiding methodological criteria. In metaphysics, one usually begins by defining a concept, and then, and only then, is one in a position to see whether there is anything in the world that satisfies that concept. I shall simply be concerned with a concept of dispositional properties that promises (so it would seem) a straightforward explanation of regularity.

3.3.1 The Dispositional / Categorical Distinction and Realism vs. Reductionism

My way of drawing the dispositional/categorical distinction is as follows:

Disposition: D is a *dispositional property* if and only if there exist a property M and conditions C such that, *necessarily*, $(\forall y)((y \text{ has } D \text{ and } y \text{ is in } C) \supset y \text{ is } M)$.⁷

Disposition explains why dispositions entail subjunctives; they entail subjunctives in virtue of the necessity in the analysans. Having stipulated the dispositional/categorical distinction, we must now ask a further question: are there any basic dispositional properties? There are two ways for a property to be dispositional (that is, to play the role of D in Disposition). First, it could satisfy

⁷ I thank Michael Tooley for proposing this type of analysis.

the right-hand side of the biconditional in virtue of its own intrinsic nature, that is, in a basic and irreducible way. Call such a property a *bare disposition*, and call the view that there are bare dispositions *bare dispositionalism*. Second, it could satisfy the right-hand side of the biconditional in virtue of features of the world extrinsic to it. Let a *categorical property* be a property that does not satisfy *D* in virtue of its own intrinsic nature. Then a categorical property *Q* could play the role of *D* in virtue of a (necessary) governing law that relates *Q* to *M* and *C* in the relevant manner. On this reductive view of dispositions *all* dispositions reduce to or supervene on other features of the world; usually, these features include (or require) laws of nature. Accordingly, reducible dispositions cannot be used to justify laws or explain regularities, since it is laws or regularities (depending on the type of reductionism) that ultimately ground dispositions.⁸ Our task is now to consider whether there is anything a priori objectionable about bare dispositional properties. If so, then one might worry that the arguments can be generalized to apply against all versions of Non-Humeanism.

It is important to note that Disposition is a (much) stronger analysis of dispositions than is popular. It can be weakened by replacing the necessity in Disposition with counterfactual dependence or by changing the analysis so that it permits finks and antidotes (see (Bird 2007) for a good example of a weaker analysis). However, there is a very good reason for using the strong analysis here. The central question we are investigating in this chapter is whether the modal elements of bare dispositions make them a priori objectionable. If objections to bare dispositions satisfying this strong analysis fail, it is unlikely that objections to bare dispositions satisfying some weaker analysis will succeed. Thus, for now our interest shall lie with the necessity present in Disposition.

⁸ I believe that this is true also for reductive possible worlds accounts of dispositions—that is, ones that rely upon something like the Stalnaker (1968) and Lewis (1973) possible worlds semantics for counterfactual conditionals.

3.3.2 Three Arguments Against Bare Dispositionalism

In this section I shall consider three arguments against bare dispositions. I argue that none provide reason to think that bare dispositionalism (or, for that matter, Non-Humeanism in general) is in any way problematic or likely to be false.

3.3.2.1 Prior, Pargetter, and Jackson

Prior, Pargetter, and Jackson (1982) endorse something like Disposition as the ground of the dispositional/categorical distinction. They argue that every disposition can be analyzed in terms of its antecedent circumstances (that is, its conditions of manifestation), manifestations, and *causal basis*, where by “causal basis” the authors mean “the property or property-complex of the object that, together with the first member of the pair—the antecedent circumstances—is the causally operative sufficient condition for the manifestation...” (251). It would seem, then, that the causal basis may be identified with the disposition itself. However, this is explicitly rejected by Prior et al. (1982). They offer instead the *distinctness thesis* (or, as MicKitrick (2003) prefers, the *non-identity thesis*):

Distinctness: Causal bases are not identical to their attendant dispositions.

The conjunction of Distinctness and Disposition entails that there are no bare dispositions (since Distinctness is effectively a reductive thesis on its own).⁹

What reasons are there for accepting Distinctness? Prior et al. (1982) offer three. The first is that the same disposition can have different causal bases in different objects. For example, the causal basis of the fragility of a wine glass differs from the causal basis of the fragility of a ceramic bowl. According to MicKitrick (2003, 358), “we do not have to say that if a disposition is ever identical to its causal basis, then it has to be identical to all of its possible causal bases.” In this case, one could claim that one of the causal bases is a bare disposition, but that the other reduces somehow. I believe that something stronger can be said in response. I see no reason to

⁹ As MicKitrick (2003, 361fn) points out, it’s not obvious that bare dispositions are the target of Prior, Pargetter, and Jackson, but their argument can be applied in this way.

accept the claim that the very same disposition can have different causal bases in different objects, because I see no reason to accept the individuation of properties required for the truth of this claim. Specifically, I see no reason to think that the fragility of the glass is the same property as the fragility of the bowl. Isn't it more plausible to suppose that we are just subsuming two distinct properties under one predicate?¹⁰ A single word can name different things. 'Jade', for example, refers to the distinct minerals *jadite* and *nephrite*. In fact, upon discovering that the fragility of the two objects is due to distinct causal bases, shouldn't we conclude that our initial judgment—that the fragility of the glass and the bowl is one and the same—is mistaken?

The second argument for Distinctness is based on the idea that dispositions can be swamped. For example, the disposition *fragility* might normally have as its causal basis the property *P*. But what if a glass with *P* also possessed some "fortifying" property *Q*, such that it was not inclined to break when struck? I believe that McKittrick's (2003, 361) response here is decisive:

As applied to a bare disposition, PPJ's suggestion amounts to the claim that an object *x* can have a disposition *D*, but simultaneously *x* can have some other properties that stop *x* from having *D*. I don't know how to make sense of this suggestion. Say you have some object *x* that is "barely fragile." You change *x*'s properties by adding some fortifying stuff to it, so that it becomes nonfragile. It is not as if, after it has become tough as nails, *x* has the bare fragility lingering inside of it. If the fragility gets "swamped," then the disposition *and* the causal basis go away. If a disposition is its causal basis, you're never going to be able to lose the disposition and keep the causal basis.

The third argument for Distinctness is based on the idea that the connection between causal basis and disposition is a contingent matter (see (Prior et al. 1982, 253-255)). Suppose that names of properties, for example, 'fragility', are rigid designators. In worlds with different laws, a glass with irregular structure might very well be tough rather than fragile. Identity facts are either necessarily true or necessarily false. Thus it is necessarily false that *having irregular structure = fragility*. Unfortunately, this argument simply begs the question against bare dispositions by

¹⁰ Michael Tooley has suggested the following point. It may be the case that there are disjunctive properties even if there are no bare dispositional properties. Thus, for example, water-solubility might be the property of having one of a collection of properties *P*, such that the laws of nature entail that anything with a property member of *P* is dissolves when placed in water. In such a case, however, it isn't at all clear to me why *P* ought to be treated as the disposition instead of its members. The latter seems more plausible as a way of characterizing the nature of the world and the things in it.

assuming that the connection between causal basis and disposition is, *in all cases*, contingent. Though I believe that the example just given is correct, this merely shows that *fragility* is not a bare disposition. But few philosophers who accept bare dispositions would claim that fragility is a bare disposition. Bare dispositions simply won't be like this.

Prior et al. (1982) do offer a response to this sort of objection. In their view, the success of this response depends on a posteriori identification of the causal basis and the disposition. But this is simply incorrect. Most proponents of bare dispositions holds that, in the case of at least some dispositions, it is a priori that the causal basis and disposition are identified. A good example is the case of an electron. Electrons are defined theoretically by their causal role. Thus, it is *a priori* that the causal basis and the disposition are one and the same. This is important because the project of explaining natural regularities relies upon dispositions that are defined similarly.¹¹

3.3.2.2 Armstrong

David Armstrong has given an objection to bare dispositions that is based on ontological economy (perhaps ontological *mystery* would be more accurate). This is the objection that, if there were bare dispositions, they would require the existence of Meinongian entities, and that this is somehow problematic. The argument comes from (Armstrong 1999). Ellis (2001, 133) offers a concise statement of the argument:

Generally speaking, things do not manifest their causal powers at all times. Indeed, some things may have causal powers that they never manifest. A causal power that had no categorical basis, and was never manifested, would have to be purely dispositional. One could not even give a relational account of it. To believe in such a property is therefore to believe in a property the essence of which involves a relation to an event that never occurs. It "points" to its manifestation, but "the manifestation does not occur." Such a property is a Meinongian property, Armstrong says, since its existence entails a relation to a non-existent event.

Why the fuss? Meinongian properties are supposedly ontologically profligate or mysterious. While I'm generally unsympathetic toward arguments from ontological economy, it is worth en-

¹¹ For a different response to this argument see MicKitrick (2003, 359-360).

tertaining this one to show where it goes wrong. With this in mind, it should be noted first that the “Meinongian” properties entailed by bare dispositions are less objectionable if they can be shown to be consistent with *actualism*, the thesis that everything actually exists. Thus, the goal is to demonstrate that the account of dispositions utilized here is consistent with actualism.

Ellis’s response, which I shall appropriate, is essentially this: causal powers don’t point to a *unique* manifestation, but rather a *whole range of possible manifestations*. Since this range of manifestations is a natural kind (that is, a universal), dispositions do not “point to” something that is ontologically mysterious; they do not entail the existence of merely possible events; they just “point to” actual properties. Consider, for example, the disposition *fragility*. To say that a glass is disposed to shatter is not to say that there exists some merely possible event of the glass’s shattering; rather, it is to say that the glass is related to the natural kind (or universal) event type, *shattering*. But this natural kind (or universal) actually exists, so Meinongian entities are not entailed by the postulate of bare dispositions. (See (Ellis 2001, 132-135) for more elaboration.) Finally, it should also be noted that Ellis’s response is consistent with Armstrong’s theory of immanent (as opposed to transcendent) universals, so long as the relata in question (the natural kinds) are somewhere instantiated.

The final point does, however, illustrate that bare dispositions have an interesting implication that Armstrong’s reductive treatment of dispositions lacks. Consider some disposition that “points to” a property that is not instantiated in the world. For example, there could be a disposition to *M* in *C*, though objects with the disposition never enter into *C*, and as a matter of chance no object in the world is *M*. If such a possibility is taken seriously (as I believe it should be), then the postulate of bare dispositions commits one to the possibility of uninstantiated, and thus transcendent, universals (since in the case described *M* is not instantiated anywhere in the actual world). Armstrong’s reductionist analysis of dispositions does not have this same implication.¹² The result is that bare dispositions may commit one to a transcendent view of properties, but not to a Meinongian view of properties. Since I believe that the former are intelligible (though the

¹² I thank Michael Tooley for explaining this point.

latter may not be), Armstrong's objection is not decisive.

3.3.2.3 Tooley

Michael Tooley has proposed an argument (in correspondence) that attempts to demonstrate that bare dispositions are impossible.¹³ The argument relies on the following analysis of the concept of an intrinsic property:

Intrinsicness: Necessarily, a property is intrinsic if and only if it is possible for a state of affairs involving that property to be the only state of affairs at a world.

The basic idea is that dispositional properties are *bare* only if they are intrinsic to the objects that possess them, but that the conjunction of bare dispositionalism and Intrinsicness entails a contradiction. For simplicity, this version of the argument considers only the class of dispositional properties that are always under their conditions of manifestation (such as the property of being disposed to conserve one's mass).

- (5) Consider the bare dispositional property *B* (that manifests under all conditions) of being disposed to preserve one's mass.
- (6) *B* is not intrinsic according to Intrinsicness (since the instantiation of *B* by any object *x* entails the existence of a later stage of *x* in which *x* has *B*).
- (7) Consider the following property that can be defined in terms of intrinsic properties plus relations:

x has the property of mass-constancy (*MC*) =_{df} if *x* has a mass = *k* at any time *t*, then for all *t** *x* exists at time *t** and has a mass = *k* at *t**.

Since the definition of *MC* is stipulative it is necessarily true. Thus, (7) entails

¹³ I thank Michael Tooley for graciously allowing me to cite his argument, and in the very simplified form in which it appears here. This simplified version is preferable for present purposes since its core idea (and, I believe, the response to it) can be presented more succinctly.

(8) Necessarily $\forall x (x \text{ has } MC \supset x \text{ conserves its mass})$.

Now recall that in Disposition, the variable D was arbitrarily chosen. Thus, it can be universally generalized to read:

Disposition: For all properties D , D is a *dispositional property* if and only if there exist a property M and conditions C such that, *necessarily*, $(\forall y)((y \text{ has } D \text{ and } y \text{ is in } C) \supset y \text{ is } M)$.

MC is a property. Since we've stipulated that there can be dispositions like B which have "vacuous" conditions of manifestation, any instantiation of MC will also be an instantiation of these vacuous conditions of manifestation. Thus the conjunction of (8) and Disposition entails that MC is a disposition. But this is problematic.

(9) Since an object can possess MC by chance (that is, not in virtue of having the bare disposition B), the possession of MC does not entail the possession of any dispositional property.

(10) Therefore, Disposition is not the correct definition of a dispositional property.

Tooley will then offer the following, revised analysis of what it is to be a dispositional property:

Disposition:* D is a *dispositional property* if and only if [D is an **intrinsic property** and there exist a property M and conditions C such that, *necessarily*, $(\forall y)((y \text{ has } D \text{ and } y \text{ is in } C) \supset y \text{ is } M)$].

Note that Disposition* does not entail that MC is a disposition; the reason, of course, is that MC is not an intrinsic property. But now bare dispositionalism faces a serious problem. The conjunction of (6) and Disposition* entails that B is *not* a disposition! But B is stipulated to be a bare disposition. Hence the postulate of bare dispositions is self-contradictory.

Though this argument appears to be powerful, I believe that the defender of bare dispositions has a successful rejoinder. We needn't accept Disposition*. There are other ways to avoid the implication that MC is a disposition. Tooley's argument implicitly relies on the assumption that if a property is not intrinsic according to Intrinsicness, then it is extrinsic in the following sense:

Extrinsicness: *Extrinsic* properties are those reducible to intrinsic properties plus relations.

But one can reasonably deny that Intrinsicness and Extrinsicness are exhaustive. There may be some properties that are neither intrinsic nor extrinsic in Tooley's sense, and bare dispositions might be properties of this sort. Thus, one could substitute the following analysis for Disposition and Disposition*:

*Disposition***: *D* is a *dispositional property* if and only if [*D* is **not an extrinsic property** and there exist a property *M* and conditions *C* such that, *necessarily*, $(\forall y)((y \text{ has } D \text{ and } y \text{ is in } C) \supset y \text{ is } M)$].

According to this analysis, *MC* is not a (nor does it entail the possession of any) disposition. The upshot of this analysis is that if there are bare dispositions, they are basic properties that are neither intrinsic nor extrinsic *in Tooley's sense*, that entail extrinsic properties, and where the entailment in question is synthetic necessary entailment.¹⁴

Of course, we want to claim that bare dispositions are intrinsic properties in some sense. If so, then an alternate analysis of intrinsicness is required. I am not completely certain what such an analysis should look like, but then I'm not completely certain what reasons would be given for thinking that something like Intrinsicness is the correct analysis in the first place. Perhaps it is intuition? If so, then Intrinsicness can be easily modified to accommodate bare dispositions, preserving the basic idea that an intrinsic property is one that does not depend ontologically on any other for its instantiation:

*Intrinsicness**: Necessarily, a property is intrinsic if and only if it is not extrinsic (in the sense of Extrinsicness) and the set consisting of a state of affairs involving that property and those states of affairs entailed by the instantiation of that property constitutes a possible world (that is, a logically/semantically consistent maximal state description).

¹⁴ Does it follow that there can be logical connections between temporally distinct states of affairs? In a word, no. It follows that there can be *necessary* connections between temporally distinct states of affairs, but these connections are *synthetic*, and hence not logical.

I'm not convinced that this is the correct analysis of intrinsic properties, but if something like it is correct (and I see no reason to think that it is not) then one can accept that bare dispositions are intrinsic properties without contradiction. Moreover, I don't see any reason why we cannot simply stipulate that this is the sense in which bare dispositions are intrinsic to their bearers.

My response to Tooley's argument can be generalized for a number of similar potential arguments against Non-Humeanism. Suppose one argues that a contradiction holds between some Non-Humean analysis and some other principle, *P*. Unless one has better empirical reasons for *P* (I can think of no empirical reasons to accept Intrinsicness or principles like it) or *P* is known a priori, the Non-Humean should just reject *P* in favor of a slightly different analysis that does not give rise to a contradiction.¹⁵ At this point, I am unaware of any suitable candidates for *P* that could harm the Non-Humean position.

3.3.3 Summing Up

To sum up, we have not discovered any serious objections to the basic idea of bare dispositions (and, a fortiori, Non-Humeanism), however unfamiliar they might seem. Over the next few chapters, I shall argue that there is very strong empirical evidence that Governing Non-Humeanism is true. Although I have argued that there are not (at this time) any a priori reasons to think that Non-Humeanism is impossible, I have not yet argued that there are good reasons for thinking that Non-Humeanism is a priori as likely as Humeanism. In the final section of this chapter, I shall just consider the a priori probabilities of Descriptive Humeanism and Governing Non-Humeanism. The a priori probability of Governing Humeanism is discussed in Chapter 4; the a priori probabilities of Descriptive Non-Humeanism and Governing Non-Humeanism are compared in Chapter 6. How should these be assigned? I do not believe that there are any a priori arguments that show that either theory is impossible, so we shall have to see what the logical interpretation of probability tells us about this matter.

¹⁵ This sort of method is employed by Helen Beebe (2000) as a response to objections to Humean Supervenience (that is, to Descriptive Humeanism).

3.4 The A Priori Probabilities of Descriptive Humeanism and Governing Non-Humeanism

Since my goal is to defend Governing Non-Humeanism, I shall ignore Governing Humeanism (until Chapter 4) and Descriptive Non-Humeanism (until Chapter 6), and instead directly compare Descriptive Humeanism and Governing Non-Humeanism here. I'm providing this comparison here (instead of in Chapter 5) because I want to motivate the idea that Governing Non-Humeanism is on a par with the other theories a priori. The matter of comparing the a priori probabilities of Descriptive Humeanism and Governing Non-Humeanism is fairly simple, at least in principle. We simply need to compare the number of state descriptions (if equiprobability of state descriptions is correct) or structure descriptions (if equiprobability of structure descriptions is correct) according to each theory. My focus shall be on the number of state descriptions, since the argument can be easily modified to show that the same result holds if instead it is equiprobability of structure descriptions that is correct. Note: Since I shall argue in Chapter 4 that Governing Humeanism is contradictory, I shall assume that Humeanism is just Descriptive Humeanism.

One final disclaimer. At the end of Chapter 5, I shall suggest that our equiprobability postulates need to be modified in order to assign probabilities to state descriptions involving probabilistic laws. I shall ignore this complication here.

So, how many state descriptions are there according to the each theory? Recall that the number of state descriptions according to Humeanism or Governing Non-Humeanism is a function of the range of values given to the parameters Σ (the number of individuals), Φ (the number of properties), and Ω (the number and type of Non-Humean connections).

Let's start with a simple example. Note that this example is merely intended to help illustrate to the reader how the interpretation of probability is to be employed; this particular example carries no weight in my argument whatsoever. Suppose that our language L is characterized as follows:

$$\Sigma = \{a,b\}$$

$$\Phi = \{F, G\}$$

$$\Omega = \{\text{necessitation}(\rightarrow)\}$$

The above assignment of parameters yields the state descriptions represented in Figures 1 and 2 (note that the relation of necessitation allows for four possible characterizations of Ω). An X next to a row refers to the fact that the state description in question is logically inconsistent; H1 refers to the first Humean state description, N1 refers to the first Non-Humean state description, and so on.

Table 3.1: Humean State Descriptions for $L: \Omega = \emptyset$

	Fa	Ga	Fb	Gb
H1	✓			
H2		✓		
H3			✓	
H4				✓
H5	✓	✓		
H6	✓		✓	
H7	✓			✓
H8		✓	✓	
H9		✓		✓
H10			✓	✓
H11	✓	✓	✓	
H12	✓	✓		✓
H13	✓		✓	✓
H14		✓	✓	✓
H15	✓	✓	✓	✓

Here is a summary of the results. The above specification of the parameters yields 15 Humean state descriptions and 19 Non-Humean state descriptions. Though some Humean state descriptions have no Non-Humean categorical equivalents (for example, H7 and H8) some Humean state descriptions have multiple Non-Humean categorical equivalents (for example H15, N8, N16, and N19 are categorical equivalents). This is the case even if there is only one type of Non-Humean connection. As the number of Non-Humean connections multiplies—one could add a relation of *exclusion*, probabilistic relations, and others—the number of Humean state descriptions without Non-Humean categorical equivalents drops, and the number of Humean state descriptions with

Table 3.2: Non-Humean State Descriptions for $L: \Omega = \{F \rightarrow G\}$

	Fa	Ga	Fb	Gb
X	✓			
N1		✓		
X			✓	
N2				✓
N3	✓	✓		
X	✓		✓	
X	✓			✓
X		✓	✓	
N4		✓		✓
N5			✓	✓
X	✓	✓	✓	
N6	✓	✓		✓
X	✓		✓	✓
N7		✓	✓	✓
N8	✓	✓	✓	✓

Table 3.3: Non-Humean State Descriptions for $L: \Omega = \{G \rightarrow F\}$

	Fa	Ga	Fb	Gb
N9	✓			
X		✓		
N10			✓	
X				✓
N11	✓	✓		
N12	✓		✓	
X	✓			✓
X		✓	✓	
X		✓		✓
N13			✓	✓
N14	✓	✓	✓	
X	✓	✓		✓
N15	✓		✓	✓
X		✓	✓	✓
N16	✓	✓	✓	✓

Table 3.4: Non-Humean State Descriptions for $L: \Omega = \{F \rightarrow G, G \rightarrow F\}$

	Fa	Ga	Fb	Gb
X	✓			
X		✓		
X			✓	
X				✓
N17	✓	✓		
X	✓		✓	
X	✓			✓
X		✓	✓	
X		✓		✓
N18			✓	✓
X	✓	✓	✓	
X	✓	✓		✓
X	✓		✓	✓
X		✓	✓	✓
N19	✓	✓	✓	✓

Non-Humean categorical equivalents increases; furthermore, the number of Non-Humean categorical equivalents that map to each Humean state description will multiply.

This is all a fancy way of saying something that we knew already: as the number of Non-Humean connections in Ω increases, the total number of Non-Humean state descriptions increases (the possibility of probabilistic connections is especially relevant here). Humeanism effectively lacks this third parameter, and thus it shouldn't be surprising that for at least some ways of describing the world it countenances fewer state descriptions than Non-Humeanism.

Ultimately, we want to know which of Humeanism and Governing Non-Humeanism contains more state descriptions (or a higher limiting relative frequency of state descriptions, in case both contain infinitely many). Calculating the number of state descriptions according to each theory would be a daunting task. This isn't a task that we must pursue, although it is important to distinguish two versions of Governing Non-Humeanism: the version according to which there *cannot* be probabilistic connections, and the version according to which there *can* be probabilistic connections—that is, probabilistic laws.

Suppose that there cannot be probabilistic laws. Then, as the number of individuals increases, the a priori probability of Humeanism increases relative to that of Governing Non-Humeanism. The reason is that although Governing Non-Humeanism will allow for multiple state descriptions that are categorically equivalent to Humean state descriptions—namely, those that have some degree of regularity—Humeanism will allow for many state descriptions that have no Non-Humean categorical equivalents. For instance, suppose that L is as specified above with the exception that Σ contains 100 individuals. In that case, the vast majority of Humean state descriptions would have no Non-Humean categorical equivalents, since the vast majority of Humean state descriptions are such as to include at least one F that is not a G and at least one G that is not an F . Thus, if there cannot be probabilistic laws then Humeanism would appear to be a priori much more likely than Governing Non-Humeanism.

The opposite is true if there can be probabilistic laws. In that case, every Humean state description will have many Non-Humean categorical equivalents. In fact, there will be infinitely

many: for each Humean state description, there will be infinitely many categorically equivalent Non-Humean state descriptions. Every specification of Σ and Φ is consistent with *every* specification of Ω that contains merely probabilistic laws. Say we have a probabilistic relation P_m , where m is the degree to which the possession of a property probabilifies the possession of some other property (the relation of nomic necessitation N is thus equivalent to the relation P_1). m can take any value between 0 and 1, and for any properties F and G , and all values of m such that $0 < m < 1$, the resulting law is consistent with every specification of Σ and Φ . Since there are infinitely many possible values, there are infinitely many state descriptions involving probabilistic laws that correspond to each Humean state description. Thus, if we restrict ourselves to cases in which Σ and Φ have finitely many members, then there are only finitely many Humean state descriptions and infinitely many Non-Humean state descriptions. Thus if probabilistic connections are intelligible, the a priori probability of Non-Humeanism (specifically, Governing Non-Humeanism) is higher than the a priori probability of Humeanism.¹⁶

But which view is correct? Can there be probabilistic laws? One option is to take an agnostic stance. Absent a proof in one direction or the other, we could average the results in the above two cases, with the result being that the a priori probability of Governing Non-Humeanism is greater than or equal to the a priori probability of Humeanism. Another option is to try to demonstrate that the possibility of probabilistic laws involves a contradiction. I do not believe that this option is likely to succeed, simply because I don't find anything strange, confusing, or otherwise mysterious about the idea of a probabilistic law. Moreover, given the present state of our most highly-respected physical theories, a rejection of probabilistic connections would seem to be an unwelcome step (though perhaps not *too* unwelcome, since there are interpretations of quantum mechanics that dispense with indeterminacy). Finally, I should note that the Humean cannot simply argue that probabilistic laws are *impossible*; she must argue that they are *unintelligible*. The reason, of course, is that the relevant probabilities are *epistemic probabilities* rather than

¹⁶ As mentioned above, this argument ignores an important complication: namely, that the equiprobability postulates in question cannot be directly applied to state descriptions involving probabilistic laws. I postpone discussion of this matter until Chapter 5.

objective chances. Thus, such an argument cannot appeal to the synthetic a priori if it is to satisfy my guiding methodological principles.

Matters are only slightly more complicated if instead it is equiprobability of structure descriptions that is correct (I shall assume the intelligibility of probabilistic laws). It would appear that Governing Non-Humeanism countenances the same number of structure descriptions as Humeanism. The reason has already be stated above. I said,

As the number of Non-Humean connections multiplies—one could add a relation of *exclusion*, probabilistic relations, and others—the number of Humean state descriptions without Non-Humean categorical equivalents drops, and the number of Humean state descriptions with Non-Humean categorical equivalents increases.

For every Humean state description H , there is at least one Non-Humean state description that is categorically equivalent to H (remember that I am assuming the intelligibility of probabilistic laws). Hence Humeanism and Non-Humeanism countenance the same number of structure descriptions. Equiprobability of Structure Descriptions then says that probabilities are assigned equally to state descriptions that are a member of each structure description. However, since we observed that there are significantly more state descriptions according to Governing Non-Humeanism than Humeanism, the suggestion will be that Governing Non-Humeanism “captures” more of the weight of each structure description. Again, the result will be that Governing Non-Humeanism is a priori more likely than Humeanism.

A worry with this explanation is that we do not know exactly how the state descriptions are distributed among the structure descriptions. Hence we should perhaps opt for a more limited conclusion. Since we don’t know the precise distribution (I’m not saying that we can’t learn it, just that as things stand at the moment we don’t know), we shouldn’t conclude that Governing Non-Humeanism does capture more weight. Absent further investigation, my suggestion is that we treat the two theories as equal.¹⁷

¹⁷ I recognize that my failure to explicitly consider this matter is far from ideal. However, my primary goal for this project is to show that, in principle, the method I have advocated in Chapter 2 may be fruitfully applied to problems in metaphysics. I believe that this can be accomplished without focusing on the somewhat painful details of the application of the method to this particular issue.

Therefore, regardless of the equiprobability postulate employed, Non-Humeanism is not assigned a significantly lower a priori probability than Humeanism. It has been suggested that Non-Humeanism is assigned a higher a priori probability than Non-Humeanism if it is assumed that probabilistic laws are intelligible—an assumption that I see no reason to reject and some positive reasons to accept.

3.5 Next Steps

This concludes my preliminary defense of Governing Non-Humeanism. I have argued (a) that Non-Humeanism is intelligible, (b) that a priori arguments against bare dispositionalism—a prominent version of Non-Humeanism—fail, and (c) that Governing Non-Humeanism is a priori just as likely as Humeanism. Chapters 4, 5, and 6 constitute my own argument for Governing Non-Humeanism. In conjunction with the arguments of this chapter, that argument demonstrates that Governing Non-Humeanism is more likely to be true (relative to the evidence of regularity) than any of the other three theories under consideration.

Chapter 4

Governing Humeanism

4.1 Introduction

Over the next three chapters I shall compare Governing Non-Humeanism to each of its three competitors. In this chapter, I consider a careful attempt to provide a Humean-consistent explanation of the necessary connection between governing laws and regularities, and I provide a general argument for thinking that no such attempt can succeed. In other words, I argue that Governing Humeanism is internally inconsistent. Since Governing *Non*-Humeanism isn't susceptible to this sort of argument, it is the superior theory.

This chapter is organized as follows. In 4.2 I discuss Michael Tooley's attempt to provide a Humean-consistent account of the entailment between governing laws and regularities, and I argue that his account faces serious problems. In 4.3 I discuss two contemporary arguments for thinking that there is no Humean-consistent account of the connection between governing law and regularity. Though I find these arguments fairly persuasive, they do have some shortcomings. In 4.4 I present a general argument for the inconsistency of Governing Humeanism that (I hope) helps to remedy the shortcomings of the other contemporary arguments. Though the argument in its present form is not decisive (I say that it is not decisive because I do not offer a complete defense of one of its crucial, but intuitively plausible, premises), it places a heavy burden on Governing Humeans to show how their view accounts for the entailment between law and regularity. The upshot is that the acceptance of Governing Laws gives us good reason to accept Non-Humeanism.

4.2 Tooley's Version of Governing Humeanism

4.2.1 Preliminaries

The most well-known version of Governing Laws (the *DTA theory*), due to Dretske (1977), Tooley (1977, 1987), and Armstrong (1983), holds that laws are states of affairs consisting of irreducible second-order external relations between universals (hence laws are atomic states of affairs distinct from regularities) that entail regularities. Consider the regularity that all *F*s are *G*s. We can explain this regularity by postulating a relation of nomic necessitation *N* between universals *F* and *G*, represented as $N(F, G)$. The crucial postulate of the theory is then

NN: For all *F* and *G*, $N(F, G)$ entails $\forall x(Fx \supset Gx)$.

N is defined as the irreducible second-order relation that satisfies *NN*. It is *NN* that explains why $N(F, G)$ explains regularities. Since $N(F, G)$ is distinct from regularities and entails regularities, $N(F, G)$ is a governing law. Our question is whether states of affairs like $N(F, G)$ are consistent with Humeanism. As we shall see later, some philosophers (notably, Bird (2005) and van Fraassen (1989, Chapter 5)) have thought that the entailment between law and regularity is something of a mystery.

Before continuing, I want to discuss *NN*'s use of the word 'entails'. Because of this word choice, it might be thought that *NN* is explicitly Humean. If so, the appropriate question is not whether states like $N(F, G)$ are consistent with Humeanism—in this case they are stipulated to be—but whether such states of affairs are possible. I prefer to interpret 'entails' more broadly, however, such that 'entails' isn't explicitly Humean, but includes the sort of synthetic necessary entailment postulated by Non-Humeans. In short, I want to interpret 'entails' in a way that is syntactically equivalent to the Humean's interpretation, but that is semantically broader so as to include synthetic necessary entailment. I think that this broader interpretation is preferable for the following reason. On the broad interpretation, it makes sense to ask what the relevant concept of entailment is: is it purely logical entailment, or is it entailment given a restricted interpretation—that is, synthetic necessary entailment? The answer to this question gives us a

straightforward answer as to whether $N(F, G)$ is consistent with Humeanism. If it can be purely logical entailment, this version of Governing Humeanism is consistent. If it cannot be, this version of Governing Humeanism is not. Moreover, the broader interpretation doesn't require us to introduce new postulates for Non-Humean governing laws, so it simplifies the dialectic.

There is another clarification worth making at this point. It is tempting to think that the existence of states of affairs like $N(F, G)$ is consistent with Humeanism on the grounds that it is true *by definition* that $N(F, G)$ entails $\forall x(Fx \supset Gx)$ —that is, on the grounds that NN is analytic. How can something that follows from a definition fail to be consistent with Humeanism? The answer is that the definition may be such that the entities capable of satisfying the relevant term(s) can only be Non-Humean entities. For example, suppose that we define *God* as the unique being in possession of the bare disposition to be obeyed by everything. Given this definition, it is analytic that God's will must be obeyed. However, worlds in which there is something that satisfies the relevant definition of God can only be Non-Humean worlds, since Humeanism precludes bare dispositions. Our worry is that the entailment posited by NN may be like the bare dispositions posited in the definition of God just given. It may be that the entailment cannot be a purely logical entailment. This matter must be investigated before we may conclude that $N(F, G)$ is a Humean-consistent state of affairs. Thus, one cannot simply conclude that NN is consistent with Humeanism on the grounds that NN is analytic. (This is why I said that a Humean consistent account of the entailment in NN must be *logical* entailment, rather than *logical or analytical* entailment.)

So, can the entailment in NN be a purely logical entailment consistent with Humeanism? Can there be a relation N if Humeanism is true? On the surface, it is not obvious how this could be so. $N(F, G)$ is stipulated to be an atomic state of affairs, one not reducible to any fact about the regularity that all F s are G s. By definition it *entails* facts about particular instances of F and G . Recall our definition of Humeanism:

Humeanism: For all propositions P , P is *necessarily true (false)* if and only if P is true (false) in

virtue of its meaning or logical form—that is, if and only if P is *analytically* true (false).

Humeanism implies that there are no synthetic necessary connections, but it looks like $N(F, G)$ might entail the existence of a synthetic necessary connection between two states of affairs.¹ Thus, there appears to be a tension between the two. Later in this chapter, I shall attempt to show that this tension gives rise to a genuine contradiction. For now, however, I want to focus on Michael Tooley's (1987) attempt to eliminate this tension by providing a Humean-consistent explanation of what laws are and how they entail regularities. I have chosen to focus on Tooley's account for the following reasons: first, it is the most carefully developed; second, it aligns most naturally with my methodological approach; third, Armstrong's (1983) attempt relies on epistemological assumptions that I have previously rejected (namely, the assumption that singular causal relations are immediately observable).

Before beginning, it is important to recognize that Tooley's account of the nature of laws comes in two steps. The first, which following (Sider 1992) I shall call *the stipulative account*, is Tooley's particular theoretical definition of governing laws. The stipulative account does not fully specify the intrinsic nature of nomic relations—it doesn't tell us *how* laws entail regularities—and thus is not explicitly Humean. Tooley's stipulative account of the relation of nomic necessitation, N , (prior to eliminating theoretical terms via the Ramsey/Carnap/Lewis method) is essentially just the account explained above.² N is irreducible, and thus not equivalent to facts about regularities; the instantiation of N by any pair of universals constitutes an atomic state of affairs; for all F and G , $N(F, G)$ is such that it entails $\forall x(Fx \supset Gx)$; Tooley also stipulates that N 's holding between two universals is a contingent matter, but I shall ignore this feature. The second, which I shall call *the speculative account*, is Tooley's attempt to show that there is a Humean-consistent way to specify the nature of nomic relations such that laws logically entail regularities. In other words, the speculative account is supposed to show that the stipulative account is consistent with

¹ I say that connection appears to be synthetic because it is not clear how the higher-order state of affairs in question is going to entail regularities of lower-order states of affairs.

² I'm focusing on Tooley's account of N instead of his general account of nomological relations for the sake of simplicity.

Humeanism.

4.2.2 Tooley's Speculative Account

I shall turn now to the matter of providing a speculative account of the intrinsic nature of nomic relations. I shall start by discussing a natural but unsuccessful speculative account before considering Tooley's.

One speculative account is based on the idea that the universals present in the state of affairs constitutive of a governing law are identical to the universals present in the regularities; hence, there is a relation of partial identity between the higher-order and lower-order states of affairs in question, and this relation entails the necessary connection. Unfortunately, this relation is insufficient. Consider NN, which says that $N(F, G)$ entails $\forall x(Fx \supset Gx)$. Here, it is easy to see that the entailment is due to the particular nature of N , and not to the mere fact that N is a relation between universals. Suppose that F and G stand in the *logical distinctness* relation: D . $D(F, G)$ does *not* entail that all F s are distinct from G s, since F and G may be coextensive even though they are distinct properties. For example, F could be the property *having a heart* and G the property *having a kidney*. Thus, the fact that N is a relation between universals does not establish the entailment between law and regularity; NN doesn't follow from the fact that $N(F, G)$ simply because N is a relation between F and G . Perhaps it will be objected that N is an external relation whereas D is internal. Consider an external relation, the *constant conjunction* relation: C . $C(F, G)$ *does* entail the regularity that all F s are G s, but it does so because C is analyzed in terms of the regularity that it entails. The main component that distinguishes N from C is that N is atomic whereas C is not. (C , of course, has no explanatory power with respect to regularities.) But $C(F, G)$ entails regularities concerning F and G precisely because it is *not* atomic; hence, we have made no progress with explaining the connection between $N(F, G)$ and the regularity $\forall x(Fx \supset Gx)$ by claiming that N is an external relation between universals. This does not suffice to establish the connection between governing law and regularity. What is it then? It must be the particular *intrinsic nature* of the external relation N itself that is responsible for the entailment between law and regularity. But

here the Humean must be very careful not to explicate the nature of the relation in such a way that it introduces (or relocates) synthetic necessities.

We turn now to Tooley's speculative account. For simplicity, I shall focus only on his account of the relation of nomic necessitation, N . Tooley's (1987, 123–129) reductive account is this:³

SPEC: $N(F, G)$ holds at world w if and only if (a) in w , there is a conjunctive universal, $F \& G$, and
(b) in w , F exists only as a part of conjunctive universal $F \& G$.

In order for this account to be a success, *SPEC* must entail *NN* while avoiding any synthetic necessities. Some elaboration and clarification is required in order to see how the account accomplishes this. In *Causation* (1987), Tooley spoke as though F was a *universal*, but this is somewhat misleading. We also need a careful explanation of what conjunctive universals are and an account of the “only as a part of” relation. What follows is a speculative clarification of *SPEC*.⁴

The relevant interpretation of *SPEC* involves drawing a distinction between *properties* and *universals* in the way specified by (1) through (5) below.

- (1) There are transcendent universals that have non-spatial, non-temporal existence, and whose existence does not logically supervene upon spatiotemporal states of affairs.
- (2) A particular can have a property only if a relevant universal is (or universals are) instantiated by the particular.

This is a Platonic rather than Aristotelian conception of universals. But this account does not entail that all properties are universals, because properties are understood as follows:

- (3) Two things a and b share a property if and only if a and b are similar in some respect.

Thus properties are conceived as “similarities” or “resemblances.” One way for two objects to resemble one another is for them to instantiate the same universals, but it is possible for some

³ Reminder: Though Tooley provides a reductive account of the necessity that binds laws to regularities, it should not be inferred that he has thereby provided a reductive account of the laws themselves.

⁴ Michael Tooley proposed this type of account in correspondence, but, as he has not endorsed the speculative extension of the account in print, it should not be taken to be representative of his own views. I thank Michael Tooley for graciously allowing me to discuss this speculative proposal here.

resemblances not to correspond to genuine universals. This is important because it allows for scenarios such as the following:

- (4) Two things, a and b , are similar in two respects F and G , and thus share two properties, F and G , though neither F nor G is a transcendent universal. Instead, there is only a *conjunctive universal* $F\&G$.

We are now in the position to say what a conjunctive universal is.

- (5) A *conjunctive universal* is a universal that is such that any two objects instantiating it are similar (resemble one another) in more than one respect. For example, we write $F\&G$ to denote the universal that is such that any objects instantiating it resemble one another in respects F and G .

The theory constituted by (1) through (5) suggests the following, revised version of SPEC, where ‘property’ and ‘conjunctive universal’ are as defined above:

*SPEC** $N(F, G)$ holds at world w if and only if (a) in w there is a conjunctive universal $F\&G$, and (b) in w , *property* F exists only as a part of conjunctive universal $F\&G$.

Now we can elaborate on the only as a part of relation. This has been the source of earlier criticism of Tooley’s account (see Sider (1992)). Sider’s primary complaint is that Tooley hasn’t carefully specified the nature of this mereological relation, and he argues that on no plausible mereological specification does Tooley’s speculative account succeed. However, we needn’t engage Sider’s critique here, because the relation need not be mereological in character. Let us analyze the only as a part of relation as follows:

- (6) A property P *exists only as a part of* universal Q at world $w =_{df}$ Q is the only universal at w such that its instantiation by multiple objects entails that those objects resemble in respect P .

Applying this analysis to *SPEC** yields the following:

$SPEC^{**}$ $N(F, G)$ holds at world w if and only if (a) in w there is a conjunctive universal $F\&G$, and (b) in w , $F\&G$ is the only universal such that its instantiation by multiple objects entails that those objects resemble in respect F .

$SPEC^{**}$ is Tooley's fully-developed account of the necessary connection between law and regularity. As required, $SPEC^{**}$ entails NN. Suppose that the world is such that there are only two universals present in that world: G and $F\&G$. Suppose further that a has property F . Then there must be some relevant universal present in a . In this world, that universal can only be $F\&G$. Accordingly, anything that has property F must (in this world) have property G , but it isn't the case that everything with property G must have property F since an object can have G by instantiating *either* universal G *or* universal $F\&G$. We shall now have to determine whether $SPEC^{**}$ entails NN *without employing any synthetic necessity*.

In the following three subsections, I shall discuss three different objections to Tooley's account. The first two objections assume that the solution just presented succeeds in its own right, but they show that the solution cannot be extended in other desirable directions. The third objection argues that the solution just presented does not succeed even for the simple case presented above—that it entails NN only if it introduces synthetic necessities, and thus that it is not a Humean-consistent theory.

4.2.3 The Objection from Probabilistic Laws

Tooley's reductive account cannot be extended to explain the connection between probabilistic law and regularity. At best, it shows that *deterministic* laws are Humean-consistent. Here is a quick explanation. The necessity that holds between law and regularity is, on this account, grounded in the necessity of identity. The reason that all F s are G s on the account provided by $SPEC^{**}$ is that a single universal is responsible for both the $Fness$ and $Gness$ of the individuals in question. In this respect, it is similar to the familiar cases of a posteriori necessities in which two *names* refer to a single property. On the standard view, it doesn't make sense to say that necessity of identity can be probabilistic; we can't say, for instance, that water = H_2O and that water is prob-

ably (but not necessarily) H_2O . Similarly, if properties F and G arise from the same universal, it isn't possible to claim that F s are *probably* (but not necessarily) G s. All things F must be G (in the world in question).

There are two potential reasons to worry about this limitation. First, it is nowadays popular to interpret quantum mechanics as implying that the world contains genuinely indeterministic laws. If this account is incapable of explaining the connection between *actual* laws and their corresponding regularities, it would fail to give us reason for thinking that Governing Humeanism is true of this world. For the record, I don't fully endorse this argument for the simple reason that I'm not convinced that quantum mechanics ought to be interpreted in the required way. It's not clear to me that the Copenhagen interpretation is superior to the Bohm interpretation on the realist interpretation of scientific theories that is required to create problems for Tooley's account.

Second, this limitation has important implications for the a priori probability of the theory. If Governing Humeanism precludes probabilistic laws then it countenances fewer possibilities than its Non-Humean competitor which simply takes the connections between law and regularity as basic and therefore does not preclude probabilistic laws. This gives us some reason to think that the a priori probability of Governing Humeanism is less than the a priori probability of Governing Non-Humeanism. (Note that this relative "reduction" in the a priori probability of Governing Humeanism won't confer any explanatory advantages over Governing Non-Humeanism.) Put in more familiar terms, other things equal we prefer a theory with fewer implications. Precluding probabilistic laws precludes lots of possibilities, and thus has lots of implications. Though I'm not convinced that any of the actual laws are probabilistic, I'd rather not rule them out.

4.2.4 The Objection from Temporally Extended Laws

Tooley's account cannot explain temporally extended laws—laws of the form *all Fs become Gs at a later time*—unless it incorporates a particular account of *structural universals*. However, Tooley's account will be unable to implement the required account of structural universals because that account involves primitive necessities. The failure of Tooley's account to explain tem-

porally extended laws is problematic for the following reasons: first, some actual regularities appear to be of this very form—consider laws about the half-life of radioactive elements; second, Tooley’s own account of the nature of causal laws set out in *Causation* (1987) holds that causal laws are temporally extended (this feature of his account is related to the familiar idea that causes are prior to their effects).

So, what are structural universals? David Lewis (1986a) explains them roughly as follows. First, they are universals: they can occur repeatedly, and so on. Second, anything that instantiates a structural universal has proper parts, and there is a necessary connection between the instantiating of the structural universal by the whole and the instantiating of other universals by the parts. We’ll say that structural universals “involve” these other universals, and we’ll say later what this involvement is in the context of specific theories of structural universals. For example,

suppose we have monadic universals *carbon* and *hydrogen*, instantiated by atoms of those elements; and a dyadic universal *bonded*, instantiated by pairs of atoms between which there is a covalent bond. (I should really be talking about momentary stages, but let’s leave time out of it for simplicity.) Then we have, for instance, a structural universal *methane*, which is instantiated by methane molecules. It involves the three previously mentioned universals as follows: necessarily, something instantiates *methane* if and only if it is divisible into five spatial parts c, h_1, h_2, h_3, h_4 such that c instantiates *carbon*, each of the h ’s instantiates *hydrogen*, and each of the c - h pairs instantiates *bonded*. (Lewis 1986a, 27)

In what sense does Tooley’s account require structural universals to explain temporally extended laws? Consider a law that says that the possession of P by individual a at t entails the possession of Q by a at $t + 1$. For Tooley’s account to work, P will have to exist only as a part of a certain structural universal. In other words, P will have to be a property (not a genuine universal) that can only be instantiated when a certain structural universal is instantiated. In this case, the structural universal will be analyzed in terms of an individual with two distinct spatiotemporal locations and an ordered pair of properties P and Q . But this requires a certain account of structural universals according to which structural universals cannot be *reduced to* their parts (where ‘parts’ is conceived broadly so as to include their structural features). Can such universals be explained in a manner consistent with Humeanism?

Lewis (1986a) discusses three accounts of structural universals. Perhaps there are others, but, like Lewis, I can't think of any. Accordingly, the Governing Humean could respond to my arguments by proposing a fourth. I'm going to suggest that the first account is compatible with the explanation given by SPEC**, but it is obviously incompatible with Humeanism. On the other hand, the second and third accounts are incompatible with SPEC**.

The first is the *magical account*. Here is how David Lewis describes this account.

On the magical conception, a structural universal has no proper parts. . . A structural universal is never simple; it involves other, simpler, universals. . . But it is mereologically atomic. The other universals it involves are not present in it as parts. Nor are the other universals set-theoretic constituents of it; it is not a set but an individual. There is no way in which it is composed of them. (Lewis 1986a, 41)

According to this account, the connection between a structural universal and the simpler universals (or properties) it involves—we might under other circumstances call these its “constituent properties”—is an unanalyzable primitive. But the connection is postulated to be a necessary one. As such, it is inconsistent with Humeanism.

Lewis elaborates:

Why *must* it be that if something instantiates *methane*, then part of it must instantiate *carbon*? According to the linguistic conception, that is built into a recursive specification of what it means to instantiate *methane*. Fair enough. According to the pictorial conception, that is because *carbon* is part of *methane*, and the whole cannot be wholly present without its part. Fair enough. But on the present conception, this necessary connection is just a brute modal fact. (Lewis 1986a, 41)

Humeans don't like brute modal facts, so we shall have to see if the linguistic or pictorial accounts can succeed where the magical account has failed.

The second is the *linguistic account*. On this account, simple universals are treated as basic predicates in a language; structural universals are treated as complex predicates defined in terms of the basic predicates in accordance with certain rules.⁵ Essentially, the linguistic conception is a reductive account of structural universals. It reduces structural universals to simple univer-

⁵ Lewis gives a much more precise explanation, but such precision is not required here.

sals plus certain logical relations (semantic rules in the relevant language). Since the account is reductive, the existence of structural universals isn't controversial, given that we have already accepted universals and the relevant logical relations (see (Lewis 1986a, 32)). However, since this is a reductive account of structural universals, it precludes an analysis of temporally extended laws in terms of SPEC**. On this view, there is no way to make sense of structural universals without taking simpler universals as basic. But if the simpler universals are basic, the properties they give rise to can't exist only as a part of structural universals; the simpler universals are simples; they exist on their own. For our temporally extended example, it holds that the structural universal incorporating property *P* is analyzed in terms of a simple (or simpler, anyway) universal *P*. Hence property *P* does not exist only as a part of the structural universal in question. On this account, universal *P* must exist on its own in order for the relevant structural universal to exist.

The third is the *pictorial account*. On this account, a structural universal is isomorphic to its instances. In other words, the structure of the universal mirrors the structure of its instances. The parts of the instances will be mirrored by the parts of the structural universal. Suppose that we have *a*'s instantiating *P* at *t* and *Q* at *t* + 1. The structural universal will involve properties *P* and *Q* standing in a certain temporal relation that mirrors the instance. This account does not fit nicely with our ordinary understanding of universals. While it is easy enough to make sense of the idea of individuals standing in temporal relations (e.g., *a* exists later than *b*), it is not at all clear that universals can stand in the sorts of temporal relations required for this account to function. It wouldn't be sufficient to say that *P* is instantiated earlier than *Q*, since that would be a reductive account of what it is for universals to stand in temporal relations and thus wouldn't do the required work. We would need to say that *P* exists (for some sense of 'exists') earlier than *Q*. What could that even mean? Further, what is to be said about *P* and *Q*? Are they universals, or merely properties (resemblances)? If they are universals, then the pictorial account faces exactly the same problems as the linguistic account: *P* doesn't exist only as a part of the structural universal in question, since universal *P* exists independently of the structural universal of which it is a part. If they are merely properties, we encounter a different problem. We have

made no progress in explaining what the relevant structure *is* and why it entails the relevant facts about particulars. If we postulate a structural universal we need to say what its structure is, lest we encounter the problems of the magical account.

The upshot of this section is that none of the above three accounts of structural universals allow Tooley's solution to be extended to the case of temporally extended laws. These three accounts may not exhaust the possibilities, but the discussion above suggests that the burden is on the Humean to show that the relevant structural universals are consistent with both Humeanism and SPEC**. A failure to do so restricts the scope of Tooley's solution significantly, since such an account is required to provide a Humean-consistent explanation of certain actual laws (such as laws about the half-life of radioactive elements) and causal laws (assuming that causal relations are cross temporal relations).

4.2.5 An Objection Concerning the Definition of Conjunctive Universals

We come now to what is in my opinion the most damaging objection to Tooley's reductive account. The concept of a conjunctive universal needs to be specified more carefully. I shall argue that conjunctive universals are capable of explaining natural regularities only if they incorporate irreducibly modal elements. Compare the following two analyses:

- (7) A *conjunctive universal* is a universal such that, *contingently*, any two objects instantiating it are similar (resemble one another) in more than one respect.
- (8) A *conjunctive universal* is a universal such that, *necessarily*, any two objects instantiating it are similar (resemble one another) in more than one respect.

Here we have a dilemma. (7) is consistent with Humeanism, but it precludes SPEC**. (8) entails SPEC**, but it precludes Humeanism. These points require some elaboration.

(7) precludes SPEC** because it entails that worlds with accidental regularities and universals are worlds with conjunctive universals. Suppose that *S* and *T* are intuitively simple (that is, non-conjunctive) universals. Suppose further that everything which has *S* also has *T*, simply as

a matter of accidental fact. According to (7), S is a conjunctive universal. But an object's having S does not entail that it has T , so laws of nature cannot be explained in terms of conjunctive universals if (7) is the correct analysis of conjunctive universals. The result is that a world with two universals G and $F\&G$ will not entail that everything with F has G ; in this world, nothing precludes $F\&G$ from giving rise only to property G . We *need* the necessary connection between $F\&G$ and properties F and G , but the account doesn't provide that unless it incorporates necessity as in (8).

Why think that the necessity in (8) is inconsistent with Humeanism? If we took conjunctive universals to be structural universals with simpler universals as their parts, it would be easy to show that the possession of a structural universal entails that any two objects instantiating it are necessarily similar in multiple respects (the relevant part of (8)); objects instantiating the structural universal would resemble in multiple ways in virtue of instantiating the simpler universals which are parts (or whatever) of the structural universal in question. Thus the necessary connection between a conjunctive universal and the properties (resemblances) to which it gives rise would be consistent with Humeanism. But Tooley cannot interpret conjunctive universals in this way, since this account entails that the simpler universals can exist without the conjunctive universals—that the simpler universals do not exist only as a part of the conjunctive universals. On this Humean-consistent account of (8), SPEC** just doesn't work. We have (8) but not SPEC**. That is, in making the necessity benign we lose SPEC**. But how else are we to explain the necessity? If we opt for a different Humean-consistent account the worry will be (again) that SPEC** doesn't follow. If we just stipulate that (8) is correct, we require an account of the necessity; it's an open question whether the necessity can be reduced. I can't think of any Humean-consistent accounts which preserve SPEC**, and for this reason I have serious doubts that SPEC** can succeed at all.

To sum up, I don't think that Tooley's speculative solution succeeds. In 4.2.5 we saw that Tooley's analysis postulates (explicitly or implicitly) modal connections. These modal connections must be analyzed in way consistent with Humeanism, and I have presented a number of

challenges for any such analysis. Even if such an analysis can meet these challenges, it will face severe limitations not faced by Non-Humean theories of governing laws, since it cannot account for probabilistic laws and it may have trouble accounting for temporally extended laws. These are reasons to prefer Governing Non-Humeanism to Governing Humeanism. In the last two sections of this chapter, I shall consider arguments for the conclusion that Governing Humeanism is self-contradictory. The result will be that we have reason to believe that there is no Humean-consistent analysis of the connection between law and regularity.

4.3 Contemporary Arguments for the Inconsistency of Governing Humeanism

In the section above I argued that Tooley's attempt to provide a Humean-consistent explanation of the necessary connection between law and regularity fails. I shall now consider arguments for the conclusion that there can be no Humean-consistent account of the necessary connection. Before presenting my own argument it will be helpful to consider its predecessors.

4.3.1 van Fraassen's Argument

Begin with two problems:

The Identification Problem: which relation between universals is the relation N ?

The Inference Problem: what information does the statement that one property necessitates another give us about what happens and what things are like?

These problems are closely related. The nature of N must be explicated in a way that provides an answer to the inference problem. In van Fraassen's words, "the relation identified as necessitation must be such as to warrant whatever we need to be able to infer from laws" (van Fraassen 1989, 96). Most critics (van Fraassen included) have focused on the Inference Problem. If it cannot be solved, the fact that $N(F, G)$ necessitates its corresponding regularity must be primitive. If it cannot be solved, Governing Humeanism is not a tenable theory.

Van Fraassen asks us to compare N with the relation of extensional inclusion, defined as follows:

Extensional Inclusion: A is extensionally included in B exactly if all instances of A are instances of B .

Clearly, A is extensionally included in B entails that all A s are B s. We now have a dilemma. By Governing Laws, the relation N cannot be the relation of extensional inclusion (since an account of laws according to which $N =$ extensional inclusion would be a descriptive account of laws). On the other hand, “how could features so distinctly different from extensional inclusion still solve the inference problem?” (van Fraassen 1989, 97).

I have no quarrel with the first horn of the dilemma, but I do not think that van Fraassen has adequately supported the second horn. His argument begins as follows:

The law as here conceived is a singular statement about universals A and B . The conclusion to be drawn from it is about another sort of things, the particulars which are instances of A and B . True, the instances are, by being instances, intimately related to the universals. This is not enough, however, to make the inference look valid to us. (van Fraassen 1989, 97)

This idea will already be familiar; it was discussed at the beginning of Section 4.2. The crucial question is whether there are other resources available to the Humean which can connect laws to regularities in the relevant manner. Van Fraassen’s argument that there are no such resources is given by way of analogy. He asks us to consider the following invalid inference (97):

1. X knows Y

therefore

2. All children of X are children of Y

The question is whether the Humean can turn it into a valid one. Van Fraassen asks us to consider the following valid inferences (98):

1. X has the same children as Y

therefore

2.

1. X has carnal knowledge only of Y

1.5. All a person's children are born of someone of whom he or she has carnal knowledge

therefore

2.

These inferences are valid, but there are problems. The first looks too trivial; an analogous inference won't give us governing laws. The second requires an extra premise, but the worry will be that the analogous premise will introduce synthetic necessity into the account of governing laws.

In support of the idea that the analogous extra premise must involve synthetic necessity, van Fraassen says the following:

That this cannot be a matter of logic, is clear from the parallel example of parents and children born to them. Perhaps relations among parents are reflected in *corresponding* relations between their children, but it will take more than logic to find the correct correspondence function! (van Fraassen 1989, 98)

The entailment asked for cannot be logical entailment, just as a fact that is purely about the parents cannot logically imply anything about the children. (Note that it would not be purely about the parents if it described them as parents, i.e. as people having children!) So the entailment cannot be a matter of logic. (van Fraassen 1989, 102)

(Note the importance of the parenthetical clause. Van Fraassen appears to be claiming that one cannot build a solution to the inference problem into the identification problem. If that is so, it is only so on the assumption that Humeanism is true. Of course, if the entailment is anything else we must reject Humeanism.)

I find this argument to be fairly persuasive. It's hard to imagine a reductive account of the necessity between law and regularity, especially after considering the failure of Tooley's attempt.

That being said, van Fraassen doesn't provide a systematic argument for thinking that there can be *no* such account. Moreover, this particular argument depends on the following analogy: parents are to second-order relations between universals as children are to regularities. This analogy is somewhat suspect even for the case at hand, and it is not at all clear that it can be extended to other accounts of governing laws that do not hold that laws are relations between universals. In sum, we want a broader argument that explicitly identifies an inconsistency in Governing Humeanism.

4.3.2 Bird's Argument

Alexander Bird (2005) has attempted to explicitly identify the contradiction in Governing Humeanism.⁶ Bird's argument considers three accounts of the necessitation between governing law and regularity.

The necessitation cannot be material implication. The resulting account would fail to satisfy Governing Laws.

The necessitation cannot be the relation of nomic necessitation. That would lead to a vicious regress of higher-order laws.

Finally, the necessitation cannot be logical entailment. Compare $N(F, G)$ with $P(F, G)$, where P is a relation of probabilification such that its holding between two universals F and G entails that there is a probability P that each instance of F will be an instance of G .

(I) $N(F, G)$ entails $\forall x(Fx \supset Gx)$

(II) $P(F, G)$ does not entail $\forall x(Fx \supset Gx)$

According to Bird, $N(F, G)$ has a modal property that $P(F, G)$ does not. Namely, it necessarily stands in a unique relation to regularities. According to Humeanism, however, there shouldn't be any states of affairs that stand in unique modal relations to regularities (other than regularities themselves), because that would be to accept a non-trivial (that is, non-reductive) modal relation.

⁶ To be precise, Bird has attempted to identify a contradiction between Armstrong's theory of laws and Armstrong's view of properties—the view that no properties are bare dispositions. I shall discuss Bird's argument as though it were an argument against Governing Humeanism.

In other words, since (I) and (II) have exactly the same logical form, their modal properties are not purely logical.

This is a creative way of immediately demonstrating the tension between Humeanism and Governing Laws. However, it relies on the idea that there are multiple basic nomological relations—in this case, N and P . Accordingly, one way for the Governing Humean to reject this argument would be to deny the existence of multiple basic nomological relations. I believe that this is an unappealing move for the Governing Humean, but it shows that Bird's argument is not the ultimate argument against Governing Humeanism. We want an argument that shows that the very idea of Governing Humeanism is self-contradictory. Bird's argument (if it succeeds) only identifies a contradiction in a slightly augmented version of Governing Humeanism.

4.4 A Modified Argument for the Inconsistency of Governing Humeanism

I shall now present a general argument for the inconsistency of Governing Humeanism. This argument is intended to show that there cannot be any Humean account of the necessary connection between governing laws and regularities.

I believe that the key to identifying the contradiction is investigating the relevant concept of *distinctness* that occurs in Governing Laws—that is, the sense in which governing laws are states of affairs distinct from regularities. So, what does it mean to say that two states of affairs are distinct? There are two obvious candidates.

$D1$: S_1 and S_2 are distinct $=_{df}$ S_1 and S_2 do not “overlap,” that is, there is no constituent individual or property of one that is a constituent of the other.

$D2$: S_1 and S_2 are distinct $=_{df}$ S_1 does not equal S_2 .

If $D1$ is correct, it follows immediately that distinct states of affairs cannot logically entail one another.⁷ Thus interpreted, there cannot be Humean-consistent governing laws. However, $D1$

⁷ Note that some Humeans will be uncomfortable with the idea of *states of affairs* standing in any entailment relations at all, even logical ones. For the sake of argument, I shall allow Humeans to postulate logical relations between states of affairs.

seems too strong. For instance, it is reasonable to say that there is a sense in which $N(F, G)$ is distinct from the regularity that all F s are G s even though both the law and the state of affairs constitutive of the regularity contain the constituents F and G . This sense of distinctness isn't captured by D1. According to D1, $N(F, G)$ and the regularity that all F s are G s are not distinct. Thus D1 doesn't capture the sense of distinctness in Governing Laws, since it doesn't countenance DTA laws as distinct from regularities. That's unacceptable.

Unlike D1, D2 does not immediately rule out the possibility of a Humean-consistent account of governing laws. However, D2 seems too weak. For instance, it holds that the following states of affairs are distinct: S_1 and $S_1 \& S_2$. It entails that a conjunctive state of affairs is distinct from each of its conjuncts. It thus allows for some descriptive/reductive accounts of laws to satisfy Governing Laws, since a law could be a conjunctive state of affairs such that one conjunct is a regularity and the other conjunct is something else. That's not acceptable either.

We need a middle ground: a definition that preserves the distinction between Governing and Descriptive Laws without holding that DTA laws are not governing laws. I think that D2 is closer to the relevant concept of distinctness. We just need to revise it so that it doesn't allow for reductive accounts of laws to satisfy Governing Laws. I don't want to argue that there aren't conjunctive states of affairs, since one plausible way of specifying regularities is as a conjunctive state of affairs satisfying certain conditions. Here, then, is an approach which allows for conjunctive states of affairs but which precludes descriptive/reductive accounts of laws from satisfying Governing Laws:

D3: States of affairs S_1 and S_2 are distinct $=_{df}$ all of the following:

- (i) $S_1 \neq S_2$
- (ii) $S_1 \neq$ any part of S_2
- (iii) $S_2 \neq$ any part of S_1

In this context, talk of parts is shorthand for the following logical concept: x is a part of $y =_{df}$ either $x = y$ or y is a conjunctive state of affairs and $x =$ one of y 's conjuncts. We are allowed

to restrict the relevant concept of parthood in this way because we are dealing with states of affairs and because we are assuming that Humeanism is true. Humeanism does not seem to be compatible with disjunctive, conditional, or negative states of affairs. They are mysterious in exactly the same way that synthetic necessity is mysterious. (More on this in a moment.)

D3 avoids the problems of D1 and D2. According to D3, DTA laws constitute governing laws, but reductive accounts of laws do not. The question now is how D3 can be employed in an argument against Governing Humeanism. Here is my proposal:

(9) The only entailment in the world consistent with Governing Humeanism is logical entailment. [from Humeanism]

(10) Governing Humeanism entails that there is some regularity. [from Governing Laws]

(11) This entailment between governing law and regularity is *not* logical entailment, since law and regularity are distinct in the sense of D3.

(12) Hence Governing Laws entails the existence of non-logical entailment in the world. [from (10) and (11)]

(13) Therefore, Governing Humeanism is self-contradictory. [from (9) and (12)]

(9) follows from Humeanism. (10) follows from Governing Laws. (12) follows from (10) and (11). (13) follows from (9) and (12). The crucial premise, obviously, is (11).

(11) says that if law and regularity are distinct in the sense of D3, one cannot logically entail the other. This sounds intuitively plausible, but let's see if it can be defended more carefully. Ideally, one would prove that (11) follows from the definition of D3—that is, that for any states of affairs p and q that are distinct in the sense of D3, p does not logically entail q and q does not logically entail p .

I'm not going to offer that sort of proof for (11). (Other critics—namely, van Fraassen (1989, Chapter 5) and Bird (2005)—haven't attempted to offer this sort of proof either.) (11) does seem

plausible, however, and I will suggest that none of the resources available to the Humean provide promising reasons to think that it is false. In other words, I shall examine all of the obvious resources for making sense of necessary connections between states of affairs and argue that none of them establish the connection between law and regularity in a manner consistent with Humeanism. This stops short of a full proof because there may be other resources available to the Humean, though I can't think of any.

We can't say that governing laws entail regularities

(14) in virtue of a semantic connection between the terms used to describe laws and regularities.

This has already been explained above, in the example in which we defined *God* as the being possessing a certain bare dispositional property. The reason is that all worlds in which there are objects satisfying the relevant definition may be Non-Humean worlds. We can describe the general problem here as follows: a merely semantic connection is something having to do with language rather than the world, but laws and regularities are features of the world.

We can't say that governing laws entail regularities because the states of affairs constitutive of regularities are such that

(15) the terms *F* and *G* used to describe states like $N(F, G)$ denote the same universal.

This is essentially Tooley's speculative proposal, and we have already seen that it is subject to serious objections. Most importantly, the idea that a single universal can give rise to multiple resemblances must be further developed. If a single universal *contingently* gives rise to multiple resemblances, this sort of account does not entail regularities. If a single universal *necessarily* gives rise to multiple resemblances, the problem of reducing or analyzing the relevant necessity has merely been moved up one level.

We can't say that governing laws entail regularities because the states of affairs constitutive of regularities are such that

- (16) governing laws are conjunctive states of affairs and the states constitutive of regularities are (some of) their conjuncts.

This account isn't consistent with Governing Laws. It holds that regularities are literally parts of the laws themselves, so laws are not distinct from regularities in the sense required by D3.

We can't say that governing laws entail regularities because the structure of the world mirrors that described by our system of logic, such that

- (17) there are *disjunctive* states of affairs, *conditional* states of affairs, and *negative* states of affairs; the conjunction of these states of affairs with regular old (atomic or conjunctive) states of affairs entail regularities.

Accepting this proposal is akin to accepting Non-Humeanism. We don't observe disjunctive or conditional or negative states of affairs in the same way that we don't observe synthetic modal connections between states of affairs. If we were to accept this explanation of the connection between governing laws and regularities we might as well accept Non-Humeanism.

I can't think of any other resources to which the Humean might appeal in order to explain why governing laws entail regularities. This doesn't mean that there aren't any; a sufficiently creative Humean might be able to find a way around these criticisms. But I hope to have shown that there are good reasons to think that (11) (and therefore (13)) is very plausible. At the very least, I believe that this argument places a heavy burden on the Governing Humean to explain the connection between law and regularity. To date, this has not been accomplished.

4.5 Conclusion

Though this may surprise the reader, I'm not seriously opposed to the idea of Governing Humeanism. I just don't know how to make sense of the claim that states of affairs distinct (in the sense of D3) from regularities entail regularities without taking that entailment as basic and unanalyzable. Every problem raised for Humeanism in this chapter has been based on the refusal to interpret the entailment between law and regularity as a synthetic necessary entailment, so

none of these problems arise for Governing Non-Humeanism. Since in the previous chapter I gave an empirically respectable account of the concept of synthetic necessity, I see no reason to burden ourselves with the strict demands of Humeanism.

What if I'm mistaken about Governing Humeanism? In that case, the arguments in the following two chapters can be reinterpreted as arguments in favor of Governing Laws instead of Governing Non-Humeanism. In the next chapter, I'll explain how observed natural regularities can be used to argue for the superiority of Governing Non-Humeanism over Descriptive Humeanism.

Chapter 5

Descriptive Humeanism

5.1 Introduction

In the preceding chapter, I argued that Governing Humeanism is inconsistent, and thus that Descriptive Humeanism is the only version of Humeanism. In this chapter, I complete the argument against Humeanism by showing that Descriptive Humeanism is very likely false in light of observed natural regularities.

My argument against Descriptive Humeanism is based on the fact that the ratio of regular state descriptions to irregular state descriptions differs according to Descriptive Humeanism and Governing Non-Humeanism. The respects in which they differ entail that the probability of regularity given Descriptive Humeanism ($P(R|H)$) is significantly less than the probability of regularity given Governing Non-Humeanism ($P(R|N)$). Bayes' theorem can then be used to show that regularity is strong evidence against Humeanism. Further, since the conclusion of Chapter 3 is that the a priori probability of Governing Non-Humeanism ($P(N)$) is greater than or equal to the a priori probability of Descriptive Humeanism ($P(H)$), Bayes' theorem entails that Governing Non-Humeanism is significantly more likely to be true than Descriptive Humeanism in light of our observation of regularity. In other words, after observing regularities we should accept Governing Non-Humeanism over Descriptive Humeanism.

This chapter is organized as follows. In 5.2 I discuss crucial assumptions and explain the relevant concept of *natural regularity*. The argument outlined above is not the first argument against Descriptive Humeanism (or related theories) based on natural regularities. In 5.3 I con-

sider earlier arguments due to Armstrong (1983) and Fales (1990). My goal is not to show that these arguments fail; rather, I merely intend to highlight their potential areas of weakness so that they may be avoided by my own argument. In 5.4 I set forth and defend my version of the argument against Descriptive Humeanism. In 5.5 I consider objections to my argument. In 5.6 I consider a big picture objection that is based on the idea that our equiprobability postulates cannot be applied to languages allowing for the possibility of probabilistic laws.

5.2 Preliminaries

5.2.1 Regularity

We have an intuitive grasp of the concept of a *natural regularity*. However, for our purposes it will help to offer a precise characterization of this concept. Obviously, universal generalizations among natural properties such as that all *Fs* are *Gs* describe regularities. Similarly, any axiom in a Ramsey/Lewis best systems analysis describes a regularity.¹ We want a definition that captures these intuitive ideas (and, of course, allows for merely probabilistic regularities). Here, then, is my stipulative definition:

Natural Regularity: A *natural regularity* is a spatiotemporally ordered set of states of affairs in which types of states of affairs are distributed in finite repeated patterns.

I take it that this is an acceptable definition, because in addition to satisfying the desiderata above it appears to be met by the regularities we actually do observe.

The regularities mentioned thus far are what I call *specific regularities*, since they concern fairly specific features of the natural world. Such regularities have been the focus of the explanatory arguments due to Armstrong and Fales (and also the focus of Michael Tooley in a manuscript on the problem of induction closely related to the project of explaining natural regularities). My argument appeals to a more general concept of regularity: the *degree* of regularity of a world (or

¹ See (Lewis 1973, 73) and (Lewis 1994) for an explanation of the best systems theory. It is important to note that one may accept the method of determining the best system in order to determine the actual laws of nature without accepting the reductive Humean ontology often associated with that method.

part of a world). Thus, for example, a *fully regular* world is one in which *all* types of states of affairs are distributed in finite repeated patterns; that is, it is one in which there exists a set of general axioms that fully specifies the distribution of all types of states of affairs at the world. What is crucial here is that it is possible to generalize from observations of particular regularities at our world to its *degree of regularity*. I shall use this general concept of degree of regularity for two principal reasons. First, its generality allows for an explanatory argument against Descriptive Humeanism (and in favor of Governing Non-Humeanism) that avoids many difficulties introduced by a focus on narrow regularities. Governing Non-Humeanism doesn't explain any particular event sequence very well at all, though some narrow versions of it may. In contrast, it *does* offer a very good explanation of the general fact that our world is highly regular. Second, our world appears to have a very high degree of regularity; perhaps it is fully regular, as no empirical observation has yet suggested otherwise. It is just as apparent to us that the world (or its observable-by-us parts, anyway) has a high degree of regularity as that it contains many specific regularities; the former is just a generalization from the latter.²

5.2.2 Review of Assumptions

Before considering the arguments, it will be helpful to review my basic assumptions. First, I assume that the arguments of Chapter 3 are correct. Thus, Governing Non-Humeanism is not only intelligible, but its a priori probability equals or exceeds that of Descriptive Humeanism. Second, I assume that the arguments of Chapter 4 are correct. Hence there is no internally consistent version of Governing Humeanism, and so all versions of Humeanism are versions of Descriptive Humeanism.³ Third, I assume that the methodological approach set forth and defended in Chapter 2 is sound. Specifically, I assume that the logical interpretation of probability offers the

² It may be helpful to note that regularity is a *distributive* property: it is a property concerned with the distribution of types of states of affairs at a world. Distributive properties vary in breadth: some are very precise (for example, every *n*th state of affairs involves property *P*), others are more general (for example, states of affairs are distributed in finite repeated patterns). The appeal to distributive properties as evidence for ontological hypotheses is not new. For example, some versions of the argument from evil appeal to the general distribution of evil states of affairs.

³ It should be noted that if this assumption turns out to be false then the argument of this chapter may be easily recast as an argument in favor of the more general thesis of Governing Laws.

correct way to determine objective epistemic probabilities. We may now move on to consider some of the previous explanatory arguments against Descriptive Humeanism.

5.3 Previous Explanatory Arguments against Descriptive Humeanism

In this section I consider two different sorts of explanatory arguments against Descriptive Humeanism.⁴ But first, a disclaimer. Armstrong accepts Governing Humeanism, and it is clear that his arguments are intended to support this thesis. However, given my purposes, I shall present an analogue of his argument to apply in favor of Governing Non-Humeanism.⁵

5.3.1 David Armstrong's Argument

Consider the following quote from David Armstrong (1988, 229):

Perhaps the regularities need no explanation? If you believe that, I say, you can believe anything.

Armstrong's point is that Descriptive Humeanism's failure to offer any explanation of natural regularities is a severe mark against it. Why does Descriptive Humeanism fail to offer any explanation of natural regularities? The reason is that Descriptive Humean laws, if there are any, *depend* on the regularities; thus the regularities cannot depend on the laws. In denying any synthetic necessities, Descriptive Humeanism denies the existence of any entities or facts (distinct from regularity itself) that could increase the likelihood of regularity. Non-Humeanism, on the other hand, does not deny such entities or facts, and so the possibility of a genuine metaphysical explanation—that is, an explanation of one phenomenon in terms of a distinct phenomenon—of regularity is preserved.

Armstrong develops his argument for this position more carefully in (Armstrong 1983, 52–59, 103–106). The first step is to show that Descriptive Humeanism cannot explain regularities.

⁴ Though not discussed in this chapter, (Foster 1982-83) is a relevant and recommended article.

⁵ (Bird 2007, 86–90) includes a different sort of explanatory argument against Descriptive Humeanism, the basic idea of which is that Humean laws don't have explanatory power with respect to any of their individual instances; however, it is sufficiently different from mine that an examination of it does not aid in the construction of my argument.

This is accomplished rather easily, as explained in the paragraph above (see also 5.4.2 of this chapter for further elaboration). The second step is to show that Descriptive Humeanism's failure to explain regularities entails inductive skepticism; that is, if Descriptive Humeanism is true, we must all be skeptics about inductive inference. The third step is to demonstrate that inductive skepticism is false. A simple application of *modus tollens* then shows that Descriptive Humeanism is false. (For Armstrong's defense of the first three steps see (Armstrong 1983, 52–59).) Armstrong includes two additional steps: an argument that his own governing theory of laws does not lead inevitably to inductive skepticism,⁶ and an inference to the best explanation that this privileged theory is correct. The idea behind these latter steps is to show that his own theory avoids the objection against Descriptive Humeanism (see (Armstrong 1983, 103–106)). These steps are required since Armstrong's theory includes the postulate of very specific laws capable of explaining the specific observed regularities. Since my argument concerns very general theories—ultimately, a set of mutually exclusive and jointly exhaustive theories—it will not require these additional steps.

Rather than explaining Armstrong's defense of these steps, I shall simply explain why I think his approach to the explanatory argument is less than ideal for those who share my methodological views. Here I have three general worries.

First, I do not think that empiricists can help themselves to Armstrong's argument against inductive skepticism. Armstrong applies a style of argument given by G.E. Moore (1925) to support the third step of his argument. However, Moore's argument cannot be adopted by the empiricist. Moore's appeal to intuition in the justification of induction is akin to appealing to synthetic a priori reasoning. The type of intuition in question is independent of experience (it must be, lest it involve a circularity) and it falls short of a logical or analytic proof. Further, no empirically respectable alternative presents itself. Whatever our justification for believing inductive inferences

⁶ Here, Armstrong's idea is this: We observe a regularity, say, that all *P*s are *Q*s; we then make an inference to the best explanation that there is a governing law of nature relating *P* and *Q* such that the instantiation of the former by an individual entails the instantiation of the latter by that same individual. When we learn that some unobserved individual *i* has *P*, we infer that *i* has *Q*.

to be reliable, it cannot be purely empirical in the sense that we just “see” the inferences to be justified. The strength of an inductive inference isn’t immediately presented to the senses. Thus, if there is an empirically respectable defense of induction it must proceed in a rather different manner.⁷

Second, the fourth step encounters two challenges, both stemming from the fact that Armstrong’s theory of laws is rather narrow. The primary challenge is that the narrowness of Armstrong’s theory severely complicates the matter of assigning a priori probabilities. Without a rigorous investigation of these matters, there is no guarantee that the explanatory power of Armstrong’s theory isn’t offset by a low a priori probability. This problem may be avoidable (Tooley’s manuscript “The Justification of Induction” includes relevant material), but even so the way to avoid it involves a very complicated line of reasoning. The secondary challenge is that the narrowness of Armstrong’s theory allows for the possibility of competing, but equally explanatory, theories. Armstrong (1983, 105) himself notes that

There do seem to be other possible explanations of the regularity of the world. Perhaps, as Berkeley thought, the regularities in things reflect no power in the things themselves, but only a particular determination of the will of God to have ordinary things (‘ideas’ for Berkeley) behave in a regular manner. If it seemed best to him, he could abrogate the so-called ‘laws of nature’ tomorrow.

It’s not initially obvious how one is to apply an inference to the best explanation in favor one theory over another when the two have equal explanatory power. Armstrong’s response to this worry is simply to admit it, in the hopes that people will see that, all things considered, his explanation is preferable to Berkeley’s. My suggestion, of course, would be to compare the a priori probabilities of the two theories, but, as I have just explained, this matter is extremely complicated.

It is worth noting that there is something a bit odd about the dialectic here. Berkeley’s account (as presented by Armstrong) offers an explanation of regularity, but not (in the quick form in which it is presented) a solution to the problem of induction. As Armstrong notes, Berkeley’s God could, if he desired, change the laws tomorrow. So Armstrong’s initial argument—the purely

⁷ For one such attempt, see Michael Tooley’s manuscript “The Justification of Induction.”

epistemological one—would seem to apply as well to Berkeley’s theory as it does to Descriptive Humeanism. This isn’t (on its own) a mark against Armstrong’s basic argument, but it does show that Armstrong’s argument isn’t solely based on the explanatory weakness of Descriptive Humeanism. This is the third problem. The basic structure of Armstrong’s argument against Descriptive Humeanism is rather different from mine. Its structure is given by *modus tollens* rather than inference to the best explanation or Bayes’ theorem. In the application of *modus tollens*, it appeals to the problem of inductive skepticism, and thus sets out to accomplish more than my argument. In doing more, it runs into additional complications. We can run an explanatory argument against Descriptive Humeanism without mention of the problem of induction. I suggest that such an argument should be run first. If it succeeds, we are then in a better dialectical position to determine which more specific version of Non-Humeanism is correct. If it fails, then stronger arguments (such as Armstrong’s) seem bound to fail as well.

The lessons to take from Armstrong’s argument are the following: first, the argument should not incorporate unnecessary epistemological assumptions, since the concept of explanation that matters to us is purely metaphysical; second, the argument should be as general as possible, striving to establish that Governing Non-Humeanism is true rather than that some more specific version of Governing Non-Humeanism is true.

5.3.2 Evan Fales’s Argument

Evan Fales (1990, Chapter 4) offers a rather different version of the explanatory argument against Descriptive Humeanism. Like Armstrong, his primary focus is on whether Non-Humeanism allows for a solution to the problem of induction. Since this problem is not of present concern, I shall simply focus on the elements of Fales’s discussion that are directly relevant to the explanation of natural regularities.

Here is Fales’s (1990, 105) summary of the argument:

We observe an event-sequence which exhibits a regular, repetitive pattern. Intuitively, and in the absence of any auxiliary information, such a sequence is re-

markable. The (*a priori*) odds are against it. Of course, the odds are equally great against any other particular sequence. But *if* (and only if) the sequence is a regular one, then there is a hypothesis, the truth of which would render the objective probability of the observed sequence very high, and the probability of any alternate sequence very low. What is this hypothesis? It is that there is a necessary connection between the constantly conjoined event-types. By Bayes' Theorem, an event or event-sequence raises the probability of a hypothesis in proportion to the degree to which that hypothesis raises the expectation of its occurrence over its prior expectation. So the existence of the regular sequence enhances the chances that this hypothesis is true.

Fales's argument has the following structure: first, it is argued that regularity constitutes evidence against Descriptive Humeanism. The idea here is to show first that the probability of that specific regularity given Descriptive Humeanism is really low; this, of course, is a simple matter, and will be familiar to anyone who has studied the problem of induction (see (Fales 1990, 92–98)). Second, it is argued that there exists a Non-Humean theory T according to which the probability of that specific regularity is very high. The conclusion is that the specific regularity constitutes strong evidence for T in the sense that it raises the probability of T significantly. The basic structure of this argument is very similar to mine. There is, however, a very important difference.

Fales asks us to consider a specific event-sequence—that is, a specific regularity—and consider hypotheses which explain it. Governing Non-Humeanism is a very general theory; *it* doesn't explain any specific event sequence well at all, though sufficiently narrow versions do. The potential problem that arises, and which also arose in the context of Armstrong's position, is this: how are we to be certain that the explanatory power of the specific Non-Humean theory T in question (that is, $P(R|T)$) isn't "offset" by a very low prior probability of T (that is, $P(T)$)?

The problem is a dialectical one. In Chapter 3, I did not attempt to calculate the *a priori* probability of Descriptive Humeanism. I just tried to show that it wasn't any higher than the *a priori* probability of Governing Non-Humeanism. But, for all we know, the *a priori* probability of Descriptive Humeanism may be much higher than the *a priori* probability of T ; that will depend, among other things, on just how narrow T is, and the more explanatory power T has, the smaller its *a priori* probability. Taking stock, here is our position according to Fales's argument:

we don't know that it is not the case that $P(H) > P(T)$; we know that $P(R|T) > P(R|H)$. We cannot arrive at any really interesting conclusion concerning $P(T|R)$ and $P(H|R)$ from this. We need more information about the a priori probabilities of T and H .

Fales provides a proposal for dealing with this problem.⁸ In essence, the goal is to show that one may be justified in assigning a sufficiently high⁹ a priori probability to the narrow theory T possessing a high degree of explanatory power with respect to the specific regularity in question. His basic suggestion is that the Humean and Non-Humean (the latter is the proponent of T) go about categorizing their theories in very different ways. Since the Humean is only concerned with the Humean ontology, the relevant "theory" is nothing more than the set of Humean state descriptions. On the other hand, the Non-Humean considers theories to be something else entirely; the relevant theories are the causal/nomological—that is, Non-Humean—structures or range of causal/nomological structures that govern the distribution of states of affairs. Thus, according to the Non-Humean, "[A theory] is a conjecture about the existence (or nonexistence) of some stable underlying structure which necessitates what has happened and what will happen." (Fales 1990, 108). Using my terminology, Fales's point is that a Non-Humean theory is nothing more than a specification of Ω . This, unfortunately, is not a thesis with which I am enamored. For one, it simplifies the method of generating Non-Humean state descriptions in a way that isn't obviously acceptable. In general, Non-Humeanism countenances multiple state descriptions for each "causal structure," that is, for each characterization of the Non-Humean connections in Ω . However, the number of state descriptions consistent with each causal structure is not constant, since it depends on the nature of the causal structure as well as the quantities of individuals and properties. Hence, this approach omits information.

Perhaps there is a way to calculate more precisely the a priori probabilities of H and T , and the conditional probabilities $P(R|T)$ and $P(R|H)$. Perhaps with the right sort of rigor, one could show that the latter ratio is higher than the former, and thus develop an argument for Non-

⁸ Fales's solution is closely related to Foster's (1982-83) solution.

⁹ That is, high enough that the explanatory gains of T aren't completely offset by its low a priori probability.

Humeanism (or, for that matter, for T). However, I don't know the best way for such an argument to proceed.¹⁰

This criticism is not intended to show that Fales's argument fails. (Actually, I am optimistic that it can be salvaged.) It is simply intended to demonstrate that the appeal to *specific* regularities as evidence introduces technical complications. These complications can be avoided by focusing on the generalization that the world is regular. How? It turns out that the very general theories of Governing Non-Humeanism and Descriptive Humeanism can be compared directly as regards their explanatory power with respect to general regularities—more precisely, with respect to the fact that the world has a high degree of regularity. Though it is by no means simple to show that the probability of regularity given Governing Non-Humeanism ($P(R|N)$) is greater than the probability of regularity given Descriptive Humeanism ($P(R|H)$), such a demonstration is possible. In my opinion, this is the path of least resistance. We have already compared the priors and found that $P(N) \geq P(H)$; if, in addition, we find that $P(R|N) > P(R|H)$, that will be enough to conclude that $P(N|R) > P(H|R)$. That, of course, is the central thesis of this chapter.

5.4 The Argument from General Regularity

The purpose of this section is to defend the general version of my explanatory argument against Descriptive Humeanism. Unlike its predecessors, it is not based on some specific regular distribution of states of affairs, but rather on the generalization that our world has a high degree of regularity.

5.4.1 In Contrast with the Previous Arguments

I can now say exactly what my argument does (and does not) accomplish. I do not defend a specific version of Governing Non-Humeanism; I simply argue that those versions expressible by positing connections in the Ω parameter have a high degree of explanatory power (relative to Descriptive Humeanism) with respect to the fact that the world contains a high degree of regu-

¹⁰ Tooley provides a related argument in "The Justification of Induction."

larity. I do not provide an analysis of laws, or of whatever it is that explains regularities. For example, I take no stand on the matter of how governing laws are to be analyzed. One could postulate higher-order relations between universals, as Dretske (1977), Tooley (1977), and Armstrong (1983) have done, with the added stipulation that the connection between law and regularity is a synthetic necessary connection. One could postulate a deity with impressive powers, where these powers are a species of bare disposition (Ellis (2001, 267) appears to attribute this view to Newton, and Armstrong (1983, 105) attributes it to Berkeley). One could postulate that it is the world itself, rather than a deity, that possesses the grand dispositions in question (see, for example, (Ellis 2001)).¹¹ All of these theories appear to be versions of Governing Non-Humeanism.

Further, I take no stand on whether laws are metaphysically contingent or metaphysically necessary. Contrary to what many philosophers have thought, the view that laws are analyzed in terms of synthetic necessary connections does not imply that laws are metaphysically necessary.¹²

For example, one could augment the Dretske/Tooley/Armstrong framework by holding that the entailment between laws and regularities is a synthetic metaphysical entailment, while still holding that the matter of which universals stand in nomic relations is a contingent matter. Similarly, one could analyze laws in terms of God's preferences, holding that it is metaphysically necessary for particular events at a world to match up with God's preferences at the world, all the while holding that God's preferences are contingent. The primary reason for these omissions is that the argument does not require a stand on issues such as these. The general thesis of Governing Non-Humeanism—the acceptance of basic synthetic necessary connections between categorical properties, represented by adding connections to parameter Ω —is all that the argument requires, because all that the argument attempts to establish is that Governing Non-Humeanism (in its most general form) is true.

¹¹ Perhaps one could postulate bare subjunctive conditionals in such a way that the conditionals entail governing laws (see, for example, (Lange 2009)), though I should mention that I find the failure of this view to respect now-popular truthmaking principles somewhat unappealing.

¹² For example, Bird (2007, 1–5, 96–97) seems to think that Non-Humeanism has this consequence.

5.4.2 A Statement of the General Argument

For simplicity of presentation, I assume that Governing Non-Humeanism and Descriptive Humeanism are mutually exclusive and jointly exhaustive. This is a reasonable assumption for the moment, since the purpose of this chapter is simply to determine which of those two theories is preferable. Also, the assumption has no effect on the overall argument for Governing Non-Humeanism, since when all is said and done (at the end of Chapter 6) the theories compared are mutually exclusive and jointly exhaustive. The general structure of the argument is as follows:

- (1) Governing Non-Humeanism is at least as likely a priori as Humeanism ($P(N) \geq P(H)$).
- (2) Governing Non-Humeanism offers a better explanation of regularity than Humeanism ($P(R|N) > P(R|H)$).
- (3) The actual world has a high degree of regularity.
- (4) Therefore, by Bayes' theorem, Governing Non-Humeanism is more likely to be true in light of our observation of regularity than Humeanism ($P(N|R) > P(H|R)$).

The relevant version of Bayes' theorem is as follows:

$$P(H|R) = \frac{P(H)P(R|H)}{P(H)P(R|H) + P(N)P(R|N)}$$

The argument is valid;¹³ if (1), (2), and (3), are true, the conclusion must be true. I defended (1) in Chapter 3. (3) is an empirical observation. That leaves (2), the defense of which is the purpose of this section.

¹³ An explanation of the argument's validity: Consider the following parallel equation for $P(N|R)$:

$$P(N|R) = \frac{P(N)P(R|N)}{P(N)P(R|N) + P(H)P(R|H)}$$

Dividing the second equation by the first yields the following:

$$\frac{P(N|R)}{P(H|R)} = \frac{P(N)P(R|N)}{P(H)P(R|H)} = \frac{P(N)}{P(H)} * \frac{P(R|N)}{P(R|H)}$$

In accordance with the methodology set out in Chapter 2, I shall defend the claim that $P(R|N) > P(R|H)$, regardless of the equiprobability postulate employed. I offer two separate arguments. Each is presented under the assumption that equiprobability of state descriptions is the correct equiprobability postulate. It is then a simple matter to demonstrate that the same arguments hold if instead equiprobability of structure descriptions is correct.

First, however, I shall explain why Descriptive Humeanism offers a poor explanation of the fact that our world is highly regular, since the ideas presented in this explanation are relevant to both of the following arguments.

In the preceding chapter, we saw that Humeanism entails that there are no governing laws: that there are no states of affairs (logically/ analytically) distinct from regularities that entail regularities. Thus, if Humeanism accepts the existence of laws they must be descriptive. Such laws may be characterized as those that satisfy the following schema:

Non-Governing Laws: x is a law of nature $=_{df}$ [some fact about regularities]

Any law fitting this schema won't provide an interesting explanation of particular matters of fact, facts about regularities included. For example, consider one of the simplest satisfactions of Non-Governing Laws:

It is a law that all F s are G s $=_{df}$ all F s are G s.

The law and the regularity are the very same thing, and no thing explains itself. The addition of Humean-consistent conditions to the right-hand side doesn't change this fact.¹⁴ More precisely, the explanatory power of a Descriptive Humean theory is inversely related to its a priori probability. It is true that the probability that all F s are G s given the Descriptive Humean law that all F s are G s is 1; however, it is also true that the a priori probability of there being a Descriptive Humean law that all F s are G s is exactly the a priori probability according to Descriptive Humeanism that all F s are G s. So the sense in which Descriptive Humean laws "explain" is not

¹⁴ Bird's (2007, 86-90) explanatory argument against Humeanism includes a substantial development of this point.

metaphysically interesting. Thus, there is no possibility of lawlike (or “top-down”) explanations of regularity given Descriptive Humeanism.

Similarly, there is no possibility of “bottom-up” explanations of regularity given Descriptive Humeanism, since Humeanism precludes the postulate of bare dispositions or any other essential connections between natural properties. Hence regularities cannot be explained by appeal to the intrinsic natures of the properties involved in regularities.

How *does* Descriptive Humeanism explain the fact that our universe is highly regular—more generally, that our universe is the way that it is? It holds that the distribution of states of affairs is a matter of chance, and nothing more. We can immediately see why previous philosophers have focused on specific regularities. Given a large sample that all *F*s are *G*s, the probability of that sample occurring by chance is incredibly small. However, something similar can be said for my broader definition of regularity: whatever the patterns involved, the probability that a given pattern will repeat by chance is incredibly small.

5.4.3 Equiprobability of State Descriptions

According to equiprobability of state descriptions, $P(R|N)$ and $P(R|H)$ are determined by the following fraction:

$$P(R|T) = \frac{\text{number of } \textit{regular} \text{ state descriptions according to theory } T}{\text{total number of state descriptions according to } T}$$

The crucial challenge is to determine how these fractions differ for the two theories. I shall provide two arguments that $P(R|N) > P(R|H)$.¹⁵

5.4.3.1 Argument 1

The first argument compares $P(R|N)$ and $P(R|H)$ by comparing the following two fractions:

$$(5) P(R|N) = \frac{\text{number of } \textit{regular} \text{ Governing Non-Humean state descriptions}}{\text{total number of Governing Non-Humean state descriptions}}$$

¹⁵ Note to the reader: The second argument is simpler and easier to understand than the first; if you wish to avoid a rather technical argument, I recommend skipping ahead to the second argument, referring back to the example in the first argument when necessary.

$$(6) P(R|H) = \frac{\text{number of regular Humean state descriptions}}{\text{total number of Humean state descriptions}}$$

The idea is to argue that (5) > (6), and one way to do so is to begin by comparing the following two fractions, where $P(R_N)$ represents either the quantity of regular state descriptions according to Governing Non-Humeanism or the a priori probability that the actual state description is a Governing Non-Humean state description (it doesn't matter which since we are assuming equiprobability of state descriptions):

$$(7) \frac{P(R_N)}{P(R_H)} = \frac{\text{number of regular Governing Non-Humean state descriptions}}{\text{number of regular Humean state descriptions}}$$

$$(8) \frac{P(I_N)}{P(I_H)} = \frac{\text{number of irregular Governing Non-Humean state descriptions}}{\text{number of irregular Humean state descriptions}}$$

I shall argue that (7) > (8). Then, using a bit of math, it follows that (5) > (6). Here is the proof:¹⁶

$$(9) \left(\frac{P(R_N)}{P(R_H)} > \frac{P(I_N)}{P(I_H)} \right) \supset \left(\frac{P(R_N)}{P(I_N)} > \frac{P(R_H)}{P(I_H)} \right).^{17}$$

$$(10) \text{ For all } a, b, c, \text{ and } d > 0, \text{ if } \frac{a}{b} > \frac{c}{d} \text{ then } \frac{a}{b+a} > \frac{c}{d+c}.^{18}$$

(11) From (9) and (10), we have:

$$\left(\frac{P(R_N)}{P(R_H)} > \frac{P(I_N)}{P(I_H)} \right) \supset \left(\frac{P(R_N)}{P(I_N)+P(R_N)} > \frac{P(R_H)}{P(I_H)+P(R_H)} \right).$$

The antecedent of (11) says that (7) > (8). The consequent of (11) says that (5) > (6)—that is, that $P(R|N) > P(R|H)$. Therefore, if (7) > (8) then (5) > (6). Accordingly, I shall now argue that (7) > (8).

The argument is based on the idea that the mapping from a Descriptive Humean state description to its Governing Non-Humean categorical equivalents is one-to-many. (Recall that, according to Governing Non-Humeanism, two state descriptions are *categorically equivalent* if and only if their assignments of individuals and properties from parameters Σ and Φ are the same and they differ only with respect to the assignment of Non-Humean connections from parameter Ω . Note also that I shall call a specification of Σ and Φ an *observable state*.) There are many

¹⁶ Earlier versions of this argument were both longer and less precise; I am indebted to Michael Tooley for suggesting this manner of setting up the argument.

¹⁷ This is a theorem; the consequent is obtained from the antecedent by cross-multiplication.

¹⁸ Proof: Suppose that $\frac{a}{b} > \frac{c}{d}$. Since b and $d > 0$, $ad > bc$. Adding ac to both sides yields $ad + ac > bc + ac$. So $a(d + c) > c(b + a)$. Since a, b, c , and $d > 0$, $(d + c)$ and $(b + a) > 0$, and thus $\frac{a}{b+a} > \frac{c}{d+c}$.

Non-Humean worlds corresponding to each highly regular observable state, since (a) not all of the regularities in a highly regular world need correspond to Non-Humean laws, and (b) there are many ways of specifying probabilistic laws that can give rise to the same regularities. In this argument, I shall ignore (b) for reasons of simplicity.¹⁹ To simplify even further, I shall assume that there is only a single basic nomic relation. In an observable state with n regularities, the number of possible Governing Non-Humean laws can range from 1 to n , and there are thus $2^n - 1$ possibilities that contain at least 1 Governing Non-Humean law, since each of the n regularities can either correspond to a Governing Non-Humean law or not, and the only case “thrown out” is the one in which no regularities correspond to a Governing Non-Humean law. This shows that, given an observable state with n regularities, $P(R_N) > P(R_H)$ by a factor of $2^n - 1$. There are $2^n - 1$ Governing Non-Humean categorically equivalent state descriptions corresponding to the single Descriptive Humean state description for the given observable state.²⁰

Of course, similar remarks can be made for the claim that $P(I_N) > P(I_H)$. Once again, the factor will be $2^n - 1$. However, it is important to remember that that factor is relativized to individual Humean state descriptions, where the value of n depends on the number of regularities in—that is, the degree of regularity of—the Humean state description in question. The higher the degree of *irregularity* of a Descriptive Humean state description, the fewer Governing Non-Humean categorical equivalents it has, since there are fewer ways of specifying the necessary connections in Ω that allow for Governing Non-Humean categorical equivalents. In other words: more regularities, more possible laws; fewer regularities, fewer possible laws. Here is an example. Suppose you have a state in which all F s are G s, but there is no regularity among P s and Q s. That observable state is consistent with its being a governing law that all F s are G s, but not with its being a law that all P s are Q s (or its being a law that all Q s are P s). For an observable state with lots of regularities, $2^n - 1$

¹⁹ The arguments are much easier to present if we assume that laws cannot be probabilistic. Then, once the arguments have been presented for the deterministic case, it is relatively straightforward to argue that they also hold for probabilistic cases.

²⁰ If we drop the assumption of a single nomic relation, the argument would be *strengthened*. In that case, for an observable state with n regularities, the number of possible Governing Non-Humean laws can range from 1 to $n * m$, where m is a constant expressing the number of basic nomic relations (basic nomological connections in Ω). For values of $m > 1$, the probability of regularity given Governing Non-Humeanism will be higher than for $m = 1$.

is high. For an observable state with few regularities, $2^n - 1$ is comparatively small. This is a very simple way of saying that $2^n - 1$ is smaller for smaller values of n . The conclusion is that there are relatively more Governing Non-Humean categorical equivalents corresponding to the Descriptive Humean state descriptions with higher degrees of regularity, and relatively fewer Governing Non-Humean categorical equivalents corresponding to the Descriptive Humean state descriptions with lower degrees of regularity. Since n corresponds to the degree of regularity at a world, Governing Non-Humeanism countenances more categorical equivalents for regular worlds, and fewer categorical equivalents for irregular worlds. That is, Governing Non-Humean state descriptions tend to be regular, whereas Humean state descriptions do not. According to Equiprobability of State Descriptions, this entails that (7) > (8).

To recap, this argument has proceeded by looking at the set of Descriptive Humean state descriptions and asking how many Governing Non-Humean categorical equivalents correspond to each Descriptive Humean state description. In each case, the answer is $2^n - 1$, where n is the number of regularities, *a number which varies from one Descriptive Humean state description to the next*. Higher degrees of regularity correspond to higher values for n ; lower degrees of regularity correspond to lower values for n . $P(R_N)$ is the number of regular Governing Non-Humean state descriptions, and it corresponds to higher values for n . $P(I_N)$ is the number of irregular Governing Non-Humean state descriptions, and it corresponds to low values for n . Hence, $P(R_N)$ is greater than $P(R_H)$ by a factor of $2^n - 1$ where n is higher, whereas $P(I_N)$ is greater than $P(I_H)$ by a factor of $2^n - 1$ where n is low. Thus, though $\frac{P(I_N)}{P(I_H)}$ may well be greater than 1, it is greater than 1 by a smaller factor than $\frac{P(R_N)}{P(R_H)}$. Therefore, (7) > (8). And, by (11), (5) > (6). Thus, $P(R|N) > P(R|H)$.

Two further points are worth noting. First, it doesn't matter where we draw the line between regular and irregular worlds, since $2^n - 1$ is monotonic (as n increases, $2^n - 1$ increases, regardless of the value of n). Thus, the basic result is that Governing Non-Humeanism predicts a higher *degree* of regularity across the board than Descriptive Humeanism. Second, we may draw a conclusion stronger than the final inequality— $P(R|N) > P(R|H)$ —indicates. Since the ratio $\frac{2^n - 1}{n}$ becomes increasingly larger as n increases, $P(R|N)$ is in fact *much* greater than $P(R|H)$. Therefore, regularity

is *strong* evidence against Descriptive Humeanism.

At this point, it may be helpful to consider an example. (Note: In order to keep the example sufficiently small that it is illuminating, a number of fairly restrictive assumptions have had to be made—for example, concerning the definition of regularity. These restrictive assumptions are *not* adopted by any of the arguments of this chapter.) Suppose that our language L is characterized as follows:

$$\Sigma = \{a,b\}$$

$$\Phi = \{F,G\}$$

$$\Omega = \{\text{necessitation}(\rightarrow)\}$$

This example allows for the state descriptions displayed in Tables 5.1 – 5.4. An X next to a row refers to the fact that the state description in question is logically inconsistent; H1 refers to the first Humean state description, N1 refers to the first Non-Humean state description, and so on.

Table 5.1: Humean State Descriptions for $L: \Omega = \emptyset$

	Fa	Ga	Fb	Gb	
H1	✓				P
H2		✓			P
H3			✓		P
H4				✓	P
H5	✓	✓			F
H6	✓		✓		P
H7	✓			✓	I
H8		✓	✓		I
H9		✓		✓	P
H10			✓	✓	F
H11	✓	✓	✓		P
H12	✓	✓		✓	P
H13	✓		✓	✓	P
H14		✓	✓	✓	P
H15	✓	✓	✓	✓	F

A quick examination tells us that, in this case, $P(N) = 19/34$ and that $P(H) = 15/34$. But what should be said about regularity? There is, I think, a very natural way of looking at regularity

Table 5.2: Non-Humean State Descriptions for $L: \Omega = \{F \rightarrow G\}$

	Fa	Ga	Fb	Gb	
X	✓				P
N1		✓			P
X			✓		P
N2				✓	P
N3	✓	✓			F
X	✓		✓		P
X	✓			✓	I
X		✓	✓		I
N4		✓		✓	P
N5			✓	✓	F
X	✓	✓	✓		P
N6	✓	✓		✓	P
X	✓		✓	✓	P
N7		✓	✓	✓	P
N8	✓	✓	✓	✓	F

Table 5.3: Non-Humean State Descriptions for $L: \Omega = \{G \rightarrow F\}$

	Fa	Ga	Fb	Gb	
N9	✓				P
X		✓			P
N10			✓		P
X				✓	P
N11	✓	✓			F
N12	✓		✓		P
X	✓			✓	I
X		✓	✓		I
X		✓		✓	P
N13			✓	✓	F
N14	✓	✓	✓		P
X	✓	✓		✓	P
N15	✓		✓	✓	P
X		✓	✓	✓	P
N16	✓	✓	✓	✓	F

Table 5.4: Non-Humean State Descriptions for $L: \Omega = \{F \rightarrow G, G \rightarrow F\}$

	Fa	Ga	Fb	Gb	
X	✓				P
X		✓			P
X			✓		P
X				✓	P
N17	✓	✓			F
X	✓		✓		P
X	✓			✓	I
X		✓	✓		I
X		✓		✓	P
N18			✓	✓	F
X	✓	✓	✓		P
X	✓	✓		✓	P
X	✓		✓	✓	P
X		✓	✓	✓	P
N19	✓	✓	✓	✓	F

in this simple example.

A world is *fully regular* $=_{df} \forall x(Fx \equiv Gx)$

A world is *partially regular* $=_{df} (\forall x(Fx \supset Gx) \vee \forall x(Gx \supset Fx) \& \sim (\forall x(Fx \equiv Gx)))$

A world is *irregular* $=_{df} \sim (\forall x(Fx \supset Gx) \vee \forall x(Gx \supset Fx))$

Note: These definitions do not satisfy my earlier definition of regularity, but they are suitable stand-ins because they allow us to examine rather small worlds. We can then look at the figures to determine which state descriptions satisfy which definitions. The raw results are represented in the rightmost column which says whether or not each state description is *fully regular* (*F*), *partially regular* (*P*), or *irregular* (*I*).

A summary of the results: Let $K(f, h)$ = the number of fully regular state descriptions according to Descriptive Humeanism, $K(p, h)$ = the number of partially regular state descriptions according to Governing Non-Humeanism, and so on.²¹

$$(12) K(f, h) = 3$$

$$(13) K(p, h) = 10$$

$$(14) K(i, h) = 2$$

$$(15) K(f, n) = 9$$

$$(16) K(p, n) = 10$$

$$(17) K(i, n) = 0$$

Suppose for the moment that we define *regularity* as *full regularity* (that is, we want to know the probabilities of a fully regular world conditional on Descriptive Humeanism and Governing Non-Humeanism). We have²²

²¹ An anonymous referee from the *Australasian Journal of Philosophy* suggested this neat manner of setting up the problem in comments on a related paper.

²² This example, in all its details, was chosen randomly, so I was surprised that the numbers worked out so neatly.

$$(18) P(R|H) = 3/15$$

$$(19) P(R|N) = 9/19$$

$$(20) P(H|R) = 1/4$$

$$(21) P(N|R) = 3/4$$

Suppose instead that we define *regularity* as *partial or full regularity*; that is, we want to know the probabilities of a world that is at least partial regular conditional on Descriptive Humeanism and Governing Non-Humeanism. We have

$$(22) P(R|H) = 13/15$$

$$(23) P(R|N) = 19/19$$

$$(24) P(H|R) = 13/32$$

$$(25) P(N|R) = 19/32$$

This example yields the desired results. It is taken to be representative of our world *only* in that it illustrates the following points: the higher the degree of regularity of a Descriptive Humean state description, the more Governing Non-Humean categorical equivalents it will have; the higher the degree of irregularity of a Descriptive Humean state description, the fewer Governing Non-Humean categorical equivalents it will have. Those state descriptions which are maximally irregular are precluded by Governing Non-Humeanism. Some partially regular state descriptions are precluded by Governing Non-Humeanism. No fully regular state descriptions are precluded by Governing Non-Humeanism. The result, which is borne out in the calculations above, is that the ratio of regular state descriptions to irregular state descriptions is greater according to Governing Non-Humeanism than according to Descriptive Humeanism.

To this point, I have not yet considered the possibility of infinitely large worlds. A different approach to the problem of determining the probabilities of regularity conditional on Descriptive Humeanism and Governing Non-Humeanism will be helpful for infinite scenarios.

5.4.3.2 Argument 2

The same conclusion—that $P(R|N)$ is much greater than $P(R|H)$ —follows if one starts with the very intuitive idea that *irregularities* are the result of a lack of governing laws—that is, a lack of Non-Humean connections in Ω . Let a *Non-Humean property** be a property such that its instantiation entails (or probabilifies) the instantiation of another property in virtue of a relevant governing law of nature. A *Humean property** is any property that is not a Non-Humean property*.

Let *Non-Humean Monism* be the view that all properties are Non-Humean properties*—that is, that all properties stand in relations to other properties as specified by Ω . *Dualism* is the view that there are both Humean properties* and Non-Humean properties*. Suppose that Non-Humean Monism is true. In this case, regularity results of necessity (that is, if there are no probabilistic connections; if there are, then regularity is just very likely, given the average of all possible probabilistic connections). Here's why. Suppose that P stands in the relation of nomic necessitation to Q . Then the following state description will be inconsistent: $Pa \ \& \ \sim Qa \ \& \ P \rightarrow Q$. The fact that $P \rightarrow Q$ guarantees the regularity that all P s are Q s. Now, if *every* property is a Non-Humean property*, every property will be like the case of P and Q , so no irregular observable state will correspond to a Governing Non-Humean observable state. The result is regularity—that is, repeated patterns. The patterns might be complex, depending on the specification of the necessary connections entailed by the Non-Humean properties*, but, given a sufficient sample, types of states of affairs must fall into a pattern (or, in the probabilistic case, are very likely to fall into a pattern). The nature of Non-Humean properties* is such that if all properties at a world are Non-Humean properties*, nothing can give rise to irregularity.

The lesson here is that the percentage of Non-Humean properties* at a world is directly related to the probability of regularity at that world. This was observed in the simple example described by the figures in the subsection above; the Non-Humean set of worlds (c) with more Non-Humean properties* (that is, with more necessary connections) implies full regularity; the Non-Humean sets of worlds (a) and (b) with fewer Non-Humean properties* do not. A higher per-

centage of Non-Humean properties* entails a higher likelihood of regularity. One who endorses Dualism can, of course, allow for irregularity, but irregularity is due to the Humean properties*, the *lack* of Non-Humean necessary connections between said properties. Suppose that there are $K = P_N + P_H$ properties, where P_N is the quantity of Non-Humean properties* and P_H is the quantity of Humean properties*. The total number of *irregular* Non-Humean state descriptions will be a function of P_H rather than K , since it is only the Humean properties* that give rise to irregularities. (In infinite worlds, matters will be spelled out in terms of limiting relative frequencies.)

Ultimately, we should like to know the a priori value of $P(R|N)$. That is, we should like to know the most likely percentage of Non-Humean properties* given Governing Non-Humeanism. In the absence of any evidence, equiprobability of state descriptions might be thought to suggest that the answer is 50%. For example, in the absence of evidence, if someone has selected a number between 1 and 100, asks you to guess the number, and will reward you based on how close you are to this number, your best bet is to select 50. Out of all the other numbers, it has the highest a priori expected payoff. Of course, one might worry here about Bertrand paradoxes; specifically, one might wonder whether there are other standpoints from which it would make sense to assign a lower or higher percentage. I can't think of any good candidates, much less any candidates that would entail that $P(R|N)$ is anywhere near as low as $P(R|H)$. For the latter reason, I don't think that this is serious problem. Given only the information that Dualism is true, regularity is much more likely than if Humean Monism is true. Given only the information that a world w contains Governing Non-Humean laws, we should expect a reasonable degree of regularity. Certainly we should expect a higher degree of regularity than for a Humean world of the same size, since, by the postulate of Non-Humeanism, $P_H < K$, but in the Humean world, $P_H = K$.

To sum up, if equiprobability of state descriptions is correct, we have good reason to believe that $P(R|N)$ is much greater than $P(R|H)$. Thus, our observation of natural regularity constitutes strong evidence in favor of Governing Non-Humeanism.

5.4.4 Equiprobability of Structure Descriptions

In accordance with the methodology set out in Chapter 2, I shall now explain how the two arguments presented above apply even if it is equiprobability of structure descriptions that is correct instead of equiprobability of state descriptions. As explained earlier, the result will be that *any* equiprobability postulate along Carnap's λ -continuum entails that $P(R|N) > P(R|H)$.

It is relatively simple to demonstrate that the same results hold as in the case of equiprobability of state descriptions. In the case of equiprobability of structure descriptions, the probability of regularity given a particular theory T is

$$P(R|T) = \frac{\sum(\text{weighted probabilities of regular state descriptions according to } T)}{\sum(\text{weighted probabilities of all state descriptions according to } T)}$$

Above, I argued that although both $P(R_N) > P(R_H)$ by a factor of $2^n - 1$ and $P(I_N) > P(I_H)$ by a factor of $2^n - 1$, n is higher in the former case than in the latter. I did so by considering the set of Descriptive Humean *state* descriptions, and then determining the number of Non-Humean categorical equivalents for each one. That number is given by the formula $2^n - 1$, where n corresponds loosely to the degree of regularity of each Humean state description. Since n is higher for worlds with higher degrees of regularity, there are relatively more Governing Non-Humean categorical equivalents for worlds with higher degrees of regularity, and thus Governing Non-Humeanism offers a better explanation of regularity than Descriptive Humeanism. The results of this procedure can be adapted for equiprobability of *structure* descriptions, the proposal under consideration here.

The procedure for calculating $P(R_H)$ under equiprobability of structure descriptions is the following:

(26) Find all *structure* descriptions.

(27) Assign equal probabilities to each structure description.

(28) For each structure description, find the *state* descriptions that are its members and distribute the probability of that structure description equally over the state descriptions

that are its members.

- (29) Find the *regular* Descriptive Humean *state* descriptions and sum their weighted probabilities to get $P(R_H)$.

The procedure for calculating $P(R_N)$ is, of course, exactly parallel, with the only difference being that ‘Governing Non-Humean’ replaces ‘Descriptive Humean’ in (29). The only major difference for this case is that the state descriptions being considered do not have the same a priori probabilities; the homogenous ones—that is, the regular ones—have higher a priori probabilities since they are assigned higher weights.

The question now is what effect this has on the original argument given in the context of equiprobability of state descriptions. In my view, the answer is that it has no (negative) effect on my argument. The argument above showed that higher degrees of regularity/homogeneity correspond to a greater difference between the number of Governing Non-Humean and Descriptive Humean state descriptions. Thus, consider some structure description S : the higher degree of homogeneity of the state descriptions members of S , the greater the proportion of Governing Non-Humean to Descriptive Humean state descriptions that are members of S ; the lower degree of homogeneity of the state descriptions members of S , the lower the proportion of Governing Non-Humean to Descriptive Humean state descriptions that are members of S . For example, consider language L (same as above) with only two predicates, F and G , and two individuals a and b . Let S^* be the structure description with probability p according to which all F s are G s and all G s are F s. There is only one Descriptive Humean state description (H15 from Figure 1 above) that is a member of S^* , but there are three Governing Non-Humean state descriptions (N8, N16, N19) that are members of S^* . Hence, the weighted probability of each state description is $p/4$, and so Governing Non-Humeanism “captures” $3/4$ the weight of p , whereas Descriptive Humeanism “captures” only $1/4$. Now let S^{**} be the structure description with probability p according to which no F s are G s and no G s are F s. There are two Descriptive Humean state descriptions (H7 and H8) members of S^{**} , but not a single Governing Non-Humean state description that is a member of S^{**} . In this

case, the entire weight of the irregular structure description belongs to Descriptive Humeanism. My second argument established exactly the same point, albeit indirectly. Instead of focusing on the relative numbers of categorical equivalent state descriptions according to the two theories, it focused directly on their ratios. But the lesson is the same. Once again, it entails that the Governing Non-Humean state descriptions will capture more weight of the more homogenous structure descriptions.

Thus, at step (28), regardless of the argument employed, the result will be that the weight of structure descriptions giving rise to regular state descriptions will tend to be distributed towards Governing Non-Humean state descriptions (since, as the earlier arguments showed, there tend to be more of them), whereas the weight of structure descriptions giving rise to irregular state descriptions will tend to be distributed towards Descriptive Humean state descriptions (since, as the earlier arguments showed, there tend to be more of them). Thus, it is evident that the arguments that $P(R|N) > P(R|H)$ proceed unaffected.

That the arguments proceed unaffected is not surprising. The general effect of moving towards equiprobability of structure descriptions is to privilege homogenous state descriptions. There are more homogenous state descriptions according to Governing Non-Humeanism than according to Descriptive Humeanism. Thus, it is to be expected that equiprobability of structure descriptions entails that $P(R|N) > P(R|H)$.

Therefore, my conclusion—that $P(R|N) > P(R|H)$ —holds regardless of the equiprobability postulate employed. Similarly, the conclusion of Chapter 3—that $P(R) \geq P(H)$ —holds regardless of the equiprobability postulate employed. And, of course, we observe pervasive natural regularities. Therefore, by Bayes' theorem, the a posteriori probability of Governing Non-Humeanism ($P(N|R)$) is greater than the a posteriori probability of Humeanism ($P(N|H)$). Thus ends my empirical argument against Humeanism.

5.5 Objections

I shall now consider five objections to the argument given in 5.4. First, one might worry that my definition of regularity is too general. The idea here is that for any infinite (or very large) distribution of states of affairs, there is a finite pattern that systematizes it. I do not think that this is correct—I'm thinking of the sequence of decimal places that constitutes any irrational number as a counterexample.²³ Moreover, I see no reason to think that a finite pattern would “emerge” from any sequence. Hence, I do not believe that this objection warrants serious consideration.

Second, one might worry that I have ignored an observation selection effect. Suppose for the moment that we live in an infinitely large Humean universe. Then we should not expect to observe irregularity, since if we were in an irregular part of the universe life would be very unlikely to have formed, or so the objection goes. Supposing that this is correct, there remain certain features of our *actual* observation of regularity that are nonetheless remarkable. If Humeanism were true, and if we were fully aware of the phenomena of observation selection effects, we would expect to live in a universe in which (a) regularity is widespread, but only *locally*, since it is only local regularity that is required to support life, and (b) our local portion of the universe is, with a very high probability, very soon going to become irregular. But we don't observe (a); instead, we observe regularity as far as observation takes us. And, though this is merely of rhetorical importance, do any of us really believe (b)?

The third objection is structurally similar to one that has been set forth against the fine-tuning argument for the existence of God.²⁴ My argument relies on the claim that $P(R|N) > P(R|H)$. But one might argue that $P(R) = 1$ on the grounds that we observe regularity.²⁵ After all, I have assumed that we know that the world has a high degree of regularity. But if $P(R) = 1$,

²³ Irrational numbers are those which cannot be expressed by a string of repeating or terminating decimals, such as π .

²⁴ See (Sober 2003) for a defense of the objection, and (Monton 2006) for criticism of the objection.

²⁵ I am ignoring the fact that regularity of observed states of affairs does not entail regularity of the whole world.

then for all T , $P(R|T) = 1$.²⁶ Thus $P(R|N) = P(R|H)$, an unacceptable result.²⁷ Unfortunately, the present proposal would make it impossible for any old observation to count as evidence against Humeanism, or as evidence against any ontology whatsoever. This is, of course, an instance of the problem of old evidence, and any solution to this problem involves adjusting the relevant probability functions. Once the relevant adjustments are made (they involve treating $P(R) < 1$), the objection loses force.²⁸ On a related note, the objection appears to conflate two distinct ways of understanding $P(R)$: the *objective epistemic a priori probability of R* and the *a posteriori probability of R given the observation that R*. Of course, it is possible for us to believe that a proposition is true while believing that the fact that it is true is unlikely. Suppose that a dealer in a game of poker repeatedly deals herself a straight flush. This outcome is a priori unlikely, but the players believe that it occurs; they observe it. The way to make sense of this is to distinguish between the two interpretations of statements of probability. Eventually, the players will believe that the dealer's good luck constitutes evidence that she is cheating, and this is despite the fact that, according to the objection being considered, the probability that she is dealt the good hands by chance given the observation that she is dealt the good hands is equal to 1. Once the distinction is made between the different interpretations of statements of probability, the objection again loses force.

The fourth objection is based on Goodman's new riddle of induction (first introduced in (Goodman 1955)). Here is the objection: Any sequence is a regularity, because predicates can be defined in such a way that any sequence is systematizable. For example, consider a world in which all green particles have positive charge and all blue particles have negative charge prior to the year 2000. After the year 2000, however, all green particles have negative charge, and all blue particles have positive charge. The connection between greenness, blueness, and charge does not appear to

²⁶ Provided that $P(e) > 0$, this is a theorem of the probability calculus. Proof: The following are theorems of the probability calculus: (i) $P(p|e) = \frac{P(p \& e)}{P(e)}$; (ii) $P(p) = 1 \supset P(p \& e) = P(e)$. Suppose that $P(p) = 1$. By (ii), $P(p \& e) = P(e)$. Substituting equivalents, $P(p|e) = \frac{P(e)}{P(e)} = 1$. By \supset -introduction, if $P(p) = 1 \supset P(p|e) = 1$. Since p and e were chosen randomly, the result may be universalized.

²⁷ I do not think that this objection is at all compelling (since it conflates two ways of interpreting statements about probability), but it seems to be a popular one and is thus worth considering.

²⁸ See (Monton 2006, 415) for a discussion of the problem of old evidence in connection with the analogous objection to the fine-tuning argument.

be easily systematizable. But wait! Let *grue* =_{df} green prior to 2000 or blue after 2000. Then this world is easily systematizable: all *grue* particles have positive charge. In this way, predicates can be defined to systematize *any* distribution of states of affairs.

This objection is easy to answer within my methodological framework. We don't get to define the predicates; our understanding of basic properties (to which the predicates correspond) is ultimately given through immediate sense experience; properties like *green* are analytically basic, and thus suitable members of Φ ; properties like *grue* are not, and are thus unsuitable. The "regularities" generated from gruesome definitions of the predicates in Φ are not genuine. The systematization involved in the above objection is mere wordplay.²⁹

The fifth objection is that there could be other empirical evidence that bears against Non-Humeanism. Although $P(N|R \text{ and only } R)$ may be greater than $P(H|R \text{ and only } R)$, it may be the case that $P(N|E) < P(H|E)$ where E is the totality of empirical evidence. Can the Humean give an empirical argument parallel to the argument of this essay that would establish that Humeanism is more likely than Non-Humeanism on other grounds? Such an argument would require evidence that bears against Non-Humeanism, such as widespread *irregularity*. But this particular bit of evidence does not obtain—at least, we do not observe it—and I cannot think of any other clear type of empirical data that could bear against Non-Humeanism. There seems to be something special about the relationship between distributive properties and this particular ontological issue, and this relationship is due to the way in which the theories are defined. Thus, I do not think that this alternative is likely to succeed.

5.6 A Big Picture Objection

The objection considered in this subsection is not an objection to the argument that $P(R|N) > P(R|H)$. Rather, it objects to the overall argument presented in 5.4.2. Thus far (not just in this chapter, but in this whole project), I have based my arguments on the simplifying assumption that our equiprobability postulates can be applied to handle state descriptions involving probabilistic

²⁹ For a related recent discussion of this issue, see (Sider 2009, 397-402).

laws. Unfortunately, there is good reason to believe that they cannot be applied in this manner; the possibility of probabilistic laws implies that our two equiprobability postulates are false (I'll explain in a moment). This gives rise to the following dilemma: on the one hand, if probabilistic laws are possible then new and controversial equiprobability postulates must be defended; on the other hand, if probabilistic laws are impossible then a further and more complicated argument is required in order to defend the thesis that $P(N|R) > P(H|R)$, since in that case my argument that $P(N) \geq P(H)$ fails. Thus, on the second horn of the dilemma I need to argue that $\frac{P(R|N)}{P(H|N)} > \frac{P(H)}{P(N)}$ even though it may be the case that $P(H) > P(N)$. The purpose of this section is to explain this dilemma and argue that the challenges of each horn can be avoided.

5.6.1 Probabilistic Laws

In Chapter 3, I appealed to the possibility of probabilistic laws in order to argue that $P(N) \geq P(H)$. Appealing to equiprobability of state descriptions, it followed immediately that there are at least as many state descriptions involving Non-Humean governing laws as Descriptive Humean state descriptions. The reasoning was as follows. Consider a probabilistic law that, for each F , there is a probability P that it will become a G . Since P can be any number between 0 and 1, we already have the result that there are infinitely many Non-Humean state descriptions. Moreover, there will be infinitely many Non-Humean state descriptions corresponding to *each* Humean state description (and thus to each observable state). This is sufficient to show that $P(N) \geq P(H)$.

The initial problem with this proposal may be seen by examining the following case:

Case 1: Let w_1 and w_2 be state descriptions of the same size (having the same number of individuals and the very same properties) such that:

- In w_1 it is a law that for each F there is a probability of .9 that it is G , and all F s are G s.
- In w_2 it is a law that for each F there is a probability of .9 that it is G , and no F s are G s.

According to equiprobability of state descriptions, w_1 and w_2 have the same a priori probability. This is obviously false. w_1 should be assigned a much higher a priori probability than w_2 . (Similar remarks can be made for equiprobability of structure descriptions.) It appears that the contents of the probabilistic laws in question have a significant bearing on the probabilities of the state descriptions in which they occur.

What is going on here? It appears that probabilistic laws introduce another “layer” of probability. The purpose of this subsection is to better articulate what such a layer must look like and to suggest that the arguments of this chapter (the arguments for the explanatory weakness of Humeanism) succeed on whatever modified interpretation of probability is required to accommodate probabilistic laws. (I stop short of a full development because I believe that the argument for the other horn of the dilemma can be countered in a very straightforward manner.)

To begin, I’d like to remind the reader that the arguments of Chapter 3 suggest that there is no reason to assign one state description a lower a priori probability than another simply on the grounds that the former contains Non-Humean governing laws and the latter does not. For this reason, (something very much like) equiprobability of state descriptions can be maintained for cases in which Humean state descriptions are compared to Governing Non-Humean state descriptions involving deterministic laws. Consider some observable state S_1 and two state descriptions satisfying that observable state: a Humean state description H_1 and a Governing Non-Humean state description GN_1 that involves only deterministic laws. Since H_1 and GN_1 differ only in that the latter includes governing laws, there is good reason to think that these two state descriptions have equal a priori probabilities. With that in mind, I believe that a fruitful place to begin is by considering how matters ought to change if we were to consider a Governing Non-Humean state description GN_2 exactly like GN_1 except that it replaces a law of the form $N(F, G)$ with a law of the form $P_{.99}(F, G)$.³⁰ It is hard to see how any justification could be given for thinking that $P(GN_2)$ is significantly lower than $P(GN_1)$, or, more generally, that probabilistic laws are a priori less likely

³⁰ I shall use $P_m(F, G)$ to represent a nomic probabilification relation such that the instantiation of F probabilifies the instantiation of G to degree m .

than deterministic laws (assuming that one agrees that they are possible).

Merits of this proposal aside, it is incomplete. Though we don't want to say that $P(GN_2)$ is significantly lower than $P(GN_1)$, we certainly don't want them to be equal. Consider GN_3 , which is exactly like GN_2 except that a single F is not a G . Unlike GN_2 , GN_3 does not map to any Governing Non-Humean state description involving only deterministic laws. Thus, looking at state descriptions like GN_1 won't give us a guide for how to assign probabilities to state descriptions like GN_3 . That said, I believe that the general strategy of assigning probabilities to state descriptions involving probabilistic laws on the basis of their relations to categorically equivalent state descriptions involving only deterministic laws is likely to succeed. Compare GN_2 and GN_3 . The relative likelihood of these two state descriptions should be given to us by the laws. (Remember that GN_2 and GN_3 have exactly the same laws.) But how?

At this point, I think that it is helpful to adopt as a heuristic a Laplacian view of the universe. On this view, the course of history begins with a set of initial conditions and laws. The initial conditions may be thought of as first-order states of affairs. Suppose you're given some initial conditions and a complete set of deterministic laws, where a complete set is such that every event and state of affairs at the world in question is law governed. There will only be one state description compossible with that Laplacian picture (call it DL_1). Now suppose that you're given the same set of initial conditions and a set of probabilistic laws. This Laplacian picture (call it PL_1) is compossible with many state descriptions. If we draw pictures to represent these views, the DL_1 is a single line extending futureward, whereas PL_1 branches in the direction of the future. The natural proposal is that the Laplacean *pictures*—the tree-like structures DL_1 and PL_1 —are equiprobable. The probability of PL_1 will then be distributed *unequally* among its various branches (where each complete "path" from the start to the end of the tree is a state description) in accordance with the content of the laws. The sort of equiprobability postulate consistent with this proposal meets all of the desiderata explained above. It doesn't artificially privilege deterministic laws over probabilistic laws, and it shows how the content of laws is relevant to the a priori probability of a given state description in a very intuitive way. Further, on this sort of postulate, the argument for the

conclusion that $P(R|N) > P(R|H)$ merely requires small formal changes.

Of course, this type of postulate will need to be specified more precisely (and there are some difficulties that arise in attempting to defend a fully specified version of such a postulate). However, the progress so far should be sufficient to demonstrate (a) that such a postulate is plausible, (b) that such a postulate is very much in the spirit of our original equiprobability postulates, and (c) that such postulates do not change the force of my arguments for Governing Non-Humeanism. In short, this suggests that the probabilistic horn of the dilemma can be countered. That said, my real reason for stopping at this point is that I believe the deterministic horn of the dilemma can be resolved much more quickly and easily.

5.6.2 Deterministic Laws

I shall now discuss the second horn of the dilemma. The problem here is that the argument that $P(N) \geq P(H)$ provided at the end of Chapter 3 depends on the assumption that probabilistic laws are possible. Without that assumption, the argument does not establish that $P(N) \geq P(H)$. In fact, most signs point to the opposite conclusion. Thus we find ourselves in the following situation. $P(N) < P(H)$, but by how much? $P(R|H) < P(R|N)$, but by how much? In order to conclude that $P(N|R) > P(H|R)$ we require an argument that $\frac{P(R|N)}{P(R|H)} > \frac{P(H)}{P(N)}$.³¹ Earlier, I criticized Evan Fales for failing to provide an adequate argument for what was essentially that same thesis. Can such an argument be provided by focusing on the concept of general regularity? I believe that it can.

I need to argue for the following thesis:

$$\frac{P(R|N)}{P(R|H)} > \frac{P(H)}{P(N)}$$

My argument begins with the claim that the number of regular Non-Humean state descriptions is greater than the number of regular Humean state descriptions. That is, that

³¹ Note that $P(N|R) > P(H|R)$ is equivalent to $P(R|N)P(N) > P(R|H)P(H)$. Dividing both sides by $P(N)P(R|H)$ gives us $\frac{P(R|N)}{P(R|H)} > \frac{P(H)}{P(N)}$.

$$(30) R_N > R_H$$

That thesis has already been established by the arguments of 5.4.3.1.³² Now divide both sides of (30) by the total number of Governing Non-Humean state descriptions (N) to get

$$(31) \frac{R_N}{N} > \frac{R_H}{N}$$

Now divide both sides of (31) by $\frac{R_H}{H}$ to get

$$(32) \frac{\frac{R_N}{N}}{\frac{R_H}{H}} > \frac{\frac{R_H}{N}}{\frac{R_H}{H}}$$

We can reduce the right-hand side of (32) to get³³

$$(33) \frac{\frac{R_N}{N}}{\frac{R_H}{H}} > \frac{H}{N}$$

Recall that under equiprobability of state descriptions

$$P(R|T) = \frac{\text{number of regular state descriptions according to theory } T}{\text{total number of state descriptions according to } T}$$

Applying this definition to the left-hand side of (33) gives us

$$(34) \frac{P(R|N)}{P(R|H)} > \frac{H}{N}$$

Further, under equiprobability of state descriptions

$$P(T) = \frac{\text{total number of state descriptions according to } T}{\text{total number of state descriptions, period}}$$

Applying this to the right-hand side of (34) yields³⁴

$$(35) \frac{P(R|N)}{P(R|H)} > \frac{P(H)}{P(N)}$$

³² Reminder: The idea is that for any observable state with n regularities, there will be only one Descriptive Humean state description corresponding to that observable state, but $2^n - 1$ Governing Non-Humean state descriptions corresponding to that observable state.

³³ As follows: $\frac{R_H}{N} / \frac{R_H}{H} = \frac{R_H}{N} * \frac{H}{R_H} = \frac{H}{N}$

³⁴ Recall that this argument is presented under the assumption that Humeanism and Non-Humeanism are mutually exclusive and jointly exhaustive. Accordingly, the denominator of the above fraction is equal to $H + N$. Thus $P(H) = \frac{H}{H+N}$ and $P(N) = \frac{N}{H+N}$. Thus $\frac{P(H)}{P(N)} = \frac{H}{H+N} / \frac{N}{H+N}$. This equals $\frac{H}{H+N} * \frac{H+N}{N}$, which is just $\frac{H}{N}$.

This is what we set out to prove at the outset. Therefore, by demonstrating that there are more regular state descriptions according to Governing Non-Humeanism than Descriptive Humeanism we can easily prove that $P(N|R) > P(H|R)$.

In sum, the explanatory weakness of Humeanism is not offset by the fact that Humeanism may have a higher a priori probability than Governing Non-Humeanism. The evidence against Humeanism is sufficiently strong that, all things considered, we may reject it even if we do not know the a priori probabilities of Humeanism and Governing Non-Humeanism.

The same argument succeeds on the structure description approach. The very same proof gets us to (33). From there, the proof is slightly more complicated. Recall that under equiprobability of structure descriptions

$$P(R|T) = \frac{\sum(\text{weighted probabilities of } \textit{regular} \text{ state descriptions according to } T)}{\sum(\text{weighted probabilities of all state descriptions according to } T)}$$

The application of this definition to (33) is less straightforward (than the application of the relevant definition from equiprobability of state descriptions). It requires the argument set forth in 5.4.4. That argument employed this definition to show that $\frac{R_N}{R_H} = \frac{P(R|N)}{P(R|H)}$. Thus, the definition above *plus* the argument of 5.4.4 gives us (34).

Further, under equiprobability of structure descriptions

$$P(T) = \frac{\sum(\text{weighted probabilities of state descriptions according to } T)}{\sum(\text{weighted probabilities of } \textit{all} \text{ state descriptions})}$$

This principle can be applied to (34) in exactly the same way as the principle used in the context of equiprobability of state descriptions to yield (35). Thus if either equiprobability of state descriptions or equiprobability of structure descriptions is correct (and thus if any equiprobability postulate on the λ -continuum is correct), the evidence against Humeanism is sufficiently strong that, all things considered, we may reject it even if we do not know the a priori probabilities of Humeanism and Governing Non-Humeanism.

Returning to the big picture, we have been considering the claim that one *cannot* show that $P(N|R) > P(H|R)$ unless there are probabilistic laws (since probabilistic laws are required to show

that $P(N) > P(H)$). As I have demonstrated, this claim is simply mistaken. Even if probabilistic laws are impossible, one can prove that $P(N|R) > P(H|R)$.

5.7 Conclusion

I have argued that, in light of the generalization that the natural world has a high degree of regularity, Descriptive Humeanism is much less likely to be true than Governing Non-Humeanism. My arguments are thoroughly empirical (though perhaps not *completely* empirical, since they assume justification for thinking that some system of logic is correct and that the logical interpretation of probability is correct). These arguments from general regularity are straightforward, so why haven't they received much attention in the literature? I think that there is a clear reason. Most philosophers working in this area have been interested either in the philosophy of science—an area focused on specific theories with great predictive power—or on the problem of induction—a problem which is most easily understood through an examination of specific cases. Philosophers with such interests would of course focus on particular regularities. However, for those of us concerned with the more general theses of Descriptive Humeanism and Governing Non-Humeanism, specific regularities aren't required; the generalization that our world has a high degree of regularity is sufficiently strong evidence in favor of Governing Non-Humeanism over Descriptive Humeanism. This strategy allows us to avoid the problem of assigning a priori probabilities to narrow hypotheses. I happen to think that this problem is neither insuperable nor inscrutable, but it just doesn't need to be engaged for our purposes. And, in this chapter, I have shown that it can be avoided while still establishing that Descriptive Humeanism is most likely false.

Chapter 6

Descriptive Non-Humeanism

6.1 Introduction

I have argued that Governing Non-Humeanism is superior to both versions of Humeanism. Now we shall turn our sights to the debate between Governing Non-Humeanism and Descriptive Non-Humeanism. The methodology should be familiar by this point. The two theories need to be compared with respect to both their a priori probabilities and their explanatory power with respect to observed natural regularities. In this chapter, I complete the argument for Governing Non-Humeanism by arguing that it is superior to Descriptive Non-Humeanism.

Last chapter, I argued that Governing Non-Humeanism has great explanatory power with respect to observed natural regularities. Can Non-Humeans achieve those same explanatory benefits without postulating governing laws? I shall argue that they cannot; Descriptive Non-Humean explanations of observed regularities merely relocate the regularities elsewhere.

I do not intend for the arguments of this chapter to be decisive; this chapter isn't supposed to be the final word on the debate between Descriptive and Governing Non-Humeanism. My primary aim here is to highlight some promising arguments that weigh against Descriptive Non-Humeanism and in favor of Governing Non-Humeanism. There are a number of reasons for the modesty of this approach. First, the matter of selecting a version of Descriptive Non-Humeanism is complicated; there are a variety of approaches in contemporary literature, all of which differ in subtle but important ways. Compare (Ellis 2001),¹ (Mumford 2004), and (Bird 2007). These im-

¹ Ellis's view is not (insofar as I can tell) straightforwardly a version of Descriptive Non-Humeanism. He holds that

portant nuances cannot be handled sufficiently in one chapter. Second, my arguments depend on a number of assumptions concerning the nature of natural properties and the nature of observation. These assumptions are controversial, and though they fit naturally with my methodological approach they do require more support than I can give them here. Third, one of my primary aims in discussing the various theories of laws is to justify the methodology of this project by giving a concrete example of how it can provide genuine epistemic reasons for metaphysical theory choice. That has already been accomplished. Accordingly, the purpose of this chapter is not to provide anything like a comprehensive discussion of Descriptive Non-Humeanism, but instead to show how my methodology can contribute to the debate between Descriptive Non-Humeanism and Governing Non-Humeanism by modestly suggesting that observed natural regularities favor the latter theory.

This chapter is organized as follows. In 6.2 I argue that the a priori probability of Descriptive Non-Humeanism is no greater than the a priori probability of Governing Non-Humeanism. In 6.3 I explain how Descriptive Non-Humeanism attempts to explain natural regularities. In 6.4 I distinguish two ways of understanding observed regularities. I argue that empirical foundationalists must understand observed regularities as relations among categorical properties, and I explain how this argument bears on contemporary debates concerning the nature of natural properties. In particular, I argue that we must reject the view according to which all natural properties are bare dispositions. In 6.5 I argue that the explanation of regularities provided in 6.3 cannot be extended to explain regularities among categorical properties, and thus that it cannot be extended to explain *observed* regularities. I conclude by contrasting the explanatory weakness of Descriptive Non-Humeanism with the explanatory power of Governing Non-Humeanism. If my arguments are sound, the result is that Governing Non-Humeanism is more likely to be true than Descriptive Non-Humeanism.

the world is of a natural kind, where the natural kind may be understood as a disposition of sorts. Depending on the precise characterization of this disposition, we may end up with an account that satisfies Governing Laws instead of Descriptive Laws. If the disposition is *bare* then it may constitute a state of affairs distinct from regularities that entails regularities. On the other hand, if it is *reducible* to those bare dispositions instantiated *in* (but not *by*) the world then Ellis's account appears to satisfy Descriptive Laws. It is not entirely clear to me which account Ellis accepts.

6.2 The A Priori Probabilities

To begin, I shall discuss the relative a priori probabilities of Governing Non-Humeanism and Descriptive Non-Humeanism. It is not terribly difficult to isolate those features of the theories relevant to their a priori probabilities, so this section will be brief.

Consider an observable state in which all F s are G s. The Governing Non-Humean can explain this regularity by positing a governing law—for instance, by positing a state of affairs $N(F, G)$. The Descriptive Non-Humean can explain this regularity by claiming that F is a bare disposition, the instantiation of which by any object brings about the instantiation of G by that same object in all conditions of manifestation. The difference between the two views is that they postulate irreducible necessity in different places. The Governing Non-Humean adds a necessary (or probabilistic) connection to Ω ; the Descriptive Non-Humean treats F as a bare disposition of sorts—a Non-Humean predicate in Φ . By giving a sufficiently creative account of how the relevant necessary connections manifest, it looks like both theories can countenance the very same observable states; more precisely, it looks like they can offer multiple categorically equivalent state descriptions for any given Descriptive Humean state description, and in fairly similar ways. Both theories can posit probabilistic necessary connections; both theories can tweak the form in which the necessary connections manifest (are they unary, binary, et cetera? are they conjunctive, disjunctive, et cetera? and so on). In sum, since the two theories have similar resources, the natural conclusion to draw is that there is a one to one mapping from Governing Non-Humean to Descriptive Non-Humean state descriptions.² For this reason, we do not yet have any reason to think that the a priori probabilities of Governing Non-Humeanism and Descriptive Non-Humeanism differ.

This result shouldn't be terribly surprising. Most Descriptive Non-Humeans (for example, Alexander Bird) believe that facts about laws, conceived as reducible relations between universals, supervene on facts about dispositions. Most Governing Non-Humeans and Governing Humeans (for example, Michael Tooley and David Armstrong) believe that facts about dispositions super-

² I'll argue against this point later in this chapter.

vene on facts about laws. This would not be the case if, given our methodology, there were obvious reasons for thinking that one theory is a priori much more likely than the other—that is, if one could accommodate lots of observable states that the other couldn't. For this reason, I shall assume that the a priori probabilities of Governing Non-Humeanism and Descriptive Non-Humeanism are the same. If we want to decide between them, we'll need to do some empirical investigation.³

6.3 How Descriptive Non-Humeanism Explains Regularities

In the last chapter, we saw that Governing Non-Humeanism does a good job of explaining natural regularities. It does so in large part because it allows for the postulate of irreducible necessary connections. So does Descriptive Non-Humeanism, so we might expect it to have similar explanatory power. I shall now explain how most Descriptive Non-Humeans have attempted to employ irreducible necessity to explain observed regularities.

The Descriptive Non-Humean cannot posit irreducible necessary connections between properties in the manner of the Governing Non-Humean. For instance, they cannot treat the necessary connections as a relation the instantiation of which constitutes an atomic state of affairs, since that would be to countenance a governing law (there are other restrictions corresponding to other ways of getting governing laws, but I shall ignore these for the moment). Instead, Descriptive Non-Humeans must build the necessity into the properties themselves. That is, they must accept that there are (something very much like) bare dispositions, and they must employ these bare dispositions to explain natural regularities without positing governing laws of nature. Recall our earlier definition of a dispositional property:

Disposition: D is a *dispositional property* if and only if there exist a property M and conditions C such that, *necessarily*, $(\forall y)((y \text{ has } D \text{ and } y \text{ is in } C) \supset y \text{ is } M)$.

Properties which satisfy this definition intrinsically—that is, in a basic and irreducible way—are bare dispositions. I shall treat the acceptance of bare dispositions understood in this way as

³ I should point out that some philosophers have thought that there are good reasons to prefer one theory to the other on the grounds that for example laws reducible to dispositions are unable to accommodate certain forms of scientific laws. See for example (Bird 2007, 211-215) for a discussion of this issue.

an essential postulate of Descriptive Non-Humeanism.⁴ This will be controversial, because (as noted in Chapter 3) Disposition is a much stronger analysis of dispositions than is popular. It can be weakened by replacing the necessity in Disposition with counterfactual dependence or by changing the analysis so that it permits finks and antidotes (see (Bird 2007) for a good example of a weaker analysis). However, there is a very good reason for using the strong analysis here. If bare dispositions satisfying the strong analysis cannot be employed to explain regularities, there will be no hope for those satisfying weaker analyses.

Now, consider some arbitrary property *D* which satisfies Disposition. Disposition entails that it is necessary that any object which possesses *D* will attain *M* under conditions *C*. This entails the regularity that all *D* in *C* are *M*. That's all there is to it! This is a very simple, and, on the surface, plausible way to account for natural regularities.

This explanation appears to make governing laws dispensable. The state of affairs consisting of a particular's possession of *D* will literally be a part of the regularity that all *D* are *M* in *C*. According to Bird (2007, 47), such generalizations are indicative of relations between universals—for example, in our case there is a relation between *D* and *M* and *C*—but these relations are not *sui generis*, not atomic, and not themselves universals. Hence the laws aren't distinct from the regularities; *all* of the properties in question are parts of the regularities, so the laws are not governing. To quote Bird, on this view “the laws spring from within the properties themselves” (2007, 2). Thus there is no need to postulate a governing law to account for the necessary connection between *D*, *M*, and *C*.

Thus it would appear that Descriptive Non-Humeanism is on a par with Governing Non-Humeanism in its ability to explain natural regularities. However, I shall argue that the explanation provided by Descriptive Non-Humeanism falls short in a very important area. Namely, though it can explain some natural regularities (as we have just seen), it cannot explain regularities among categorical properties. This result is quite damning, since the regularities given in

⁴ In Chapter 3, I considered a slightly modified version of Disposition; this modification is ignored here for the sake of simplicity, but this omission has no effect on the arguments of this chapter.

experience—the *observed* regularities, the regularities most in need of explanation—are regularities among categorical properties. Accordingly, my argument against Descriptive Non-Humeanism has two parts. In the next section, I argue that observed regularities constitute regularities among categorical properties. In the section that follows, I argue that Descriptive Non-Humeanism is unable to provide an explanation of these regularities.

6.4 The Nature of Observed Regularities

I shall begin my argument against Descriptive Non-Humeanism by arguing that (at least some) observed regularities constitute regularities among categorical properties.

At this point, we must ask a very important question: What is the nature of properties given in experience? No doubt many philosophers will argue that the properties given in experience (supposing, of course, that some properties *are* given in experience—a reasonable supposition given the epistemological assumptions of this project) are obviously categorical properties. Though there has been little direct discussion of this issue in recent literature, there has been much discussion over a closely related issue: whether all natural properties are categorical (*categorical monism*), all natural properties are bare dispositions (*dispositional monism*), or there are natural properties of both types (*dualism*). For the purposes of this chapter we can ignore categorical monism. The relevant question, accordingly, is whether there are some categorical properties. If there are not, then obviously the natural properties given in experience cannot be categorical properties. Thus the debate over the types of natural properties is relevant to the matter of whether the properties given in experience are dispositional or categorical.

It would go beyond the scope of this chapter to provide a full assessment of the debate over the types of natural properties. Instead, I shall introduce and defend an argument for the conclusion that (at least some) observed regularities constitute relations among categorical properties only—that is, that when we observe regularities we (at least sometimes) observe regularities that hold among categorical properties. Once this argument has been introduced, I shall explain how it bears on two of the most important arguments in the debate over whether there are any

categorical properties. Ultimately, I believe that the matter hinges on one's underlying methodological assumptions. In this case, empirical foundationalism privileges the view that at least some properties—namely, the properties given in experience—are categorical.

6.4.1 The Argument From Experience

I shall now argue that observed regularities constitute regularities among categorical properties only.⁵ We are interested in the metaphysics of laws of nature in the first place because our experiences of the world are systematizable into relatively simple patterns (patterns that are a priori unlikely). In other words, we care about laws because we observe regularities. Thus it is very important to determine the nature of observed regularities. Are they regularities among categorical properties, dispositional properties, or both?

In order to answer this question, I need to introduce a few other assumptions about the nature of properties. First, I assume that natural properties are *sparse*. This is to say that not every predicate or concept corresponds to a genuine property. For instance, one could accept this view by holding that mass, charge, and the other fundamental properties of physics are the only (metaphysically) *real* properties. Second, I shall assume that real properties are universals. (We could probably do without the second assumption, but it allows for a simple presentation of the competing explanations of natural regularities.) Third, I shall assume that the instantiation of the same property by two individuals implies that the two individuals resemble one another in at least one respect. This probably goes without saying, but by 'resemble' I don't mean 'visually resemble' nor do I mean 'resemble in some way detectable by the senses'. This brings us to two very important questions: In which respects do the instantiations of a categorical property resemble one another? In which respects do the instantiations of a bare disposition resemble one another? The central idea behind my argument is that categorical properties and dispositional properties differ in these respects.

It seems obvious that there is a genuine distinction between *modal* and *non-modal* resem-

⁵ This argument is related to Swinburne's (1980) epistemological regress argument against dispositional monism.

blance relations. Consider Disposition. By stipulation, at least one respect (and, in my view, the only respect) in which two instantiations of a bare disposition resemble one another is a modal respect; the resemblance cannot be grasped without possessing the antecedent modal concepts appearing in the analysis of Disposition. There is simply no other respect in which two instantiations of a bare disposition must resemble one another. If they didn't resemble in the relevant modal respect then the relevant property could not be a bare disposition. In contrast, it is clear that some objects resemble without resembling in modal respects. Consider two instantiations of the property *qualitative squareness* in different regions of one's visual field. These property instantiations resemble, and the resemblance can be grasped without any antecedent modal concepts. So there are both modal and non-modal resemblances. The upshot is this. Since we have good reason to accept the distinction between modal and non-modal resemblances, and since by definition the only respects in which instantiations of a bare disposition must resemble are modal, it follows that we have the concept of categorical properties. Further, since some resemblances are (a) given in experience, and (b) non-modal, it follows that there are some categorical properties.

We can say something stronger: that modal resemblance relations cannot be given in experience at all; that the respect in which two instantiations of a bare disposition resemble one another is unobservable. Consider the following argument:

- (1) Two instantiations of the same bare disposition resemble in a modal respect.
- (2) We cannot be acquainted (through observation) with the modal respects in which the instantiations of a given property resemble one another.
- (3) Therefore, bare dispositions cannot be given in experience.
- (4) Since we do have experiences, some properties are given in experience.
- (5) Therefore, the properties given in experience are not dispositional properties.

Since, by definition, any property that is not a bare disposition is a categorical property, it follows that some categorical properties are observable. Since observed regularities are observable, it

follows that observed regularities are regularities among categorical properties.

I have already provided a defense of (1) (above—it simply follows from the definition of a bare disposition). *Pace* eliminativists about the mental, (4) is as obvious as anything is obvious. This leaves (2), which is widely accepted. Consider: we can be acquainted with *saltiness* easily enough; we cannot (it would seem) be acquainted with *water-solubility*. One way to support (2) is to appeal to examples such as this, but there is a better argument for (2). Begin with Hume's argument for the unobservability of causal relations. In Chapter 2 I explained that this argument may be extended to any modal facts about the world; they just aren't the sorts of facts that can be given in experience. The truth of counterfactuals isn't given in experience; causal relations aren't given in experience; governing laws of nature aren't given in experience. Why not? We experience that which *is*, but not that which merely *could be*. We taste salt; we observe crystalline structure; we observe dissolving in water; in short, we observe the regularity that all salt dissolves in water; we do *not* observe water-solubility. We know what it's like to observe *redness* and *squareness*. But what could it possibly be like to observe a property *being such that, necessarily, dissolves when placed in water*? True, this isn't exactly an argument that bare dispositions *cannot* be given in experience, but it strongly suggests that dispositional properties are not given in experience. At the very least, it places the following burden on those who believe that dispositions can be given in experience: they need to tell us how it is that such properties are given in experience and what it is like to experience them. If they can't, we have no reason to think that bare dispositional properties are given in experience, and thus we have no reason to think that observed regularities constitute regularities among bare dispositional properties. (Note that it won't be of any help to say that dispositional properties are not *given* in experience but that they *give rise to* experiences. The categorical nature of our experiences must be accounted for, and this road leads to both dualism and the view that observed regularities are relations among categorical properties.)

Of course, it is open to the Descriptive Non-Humean to deny (2). Indeed, I believe that this is the most promising strategy for the Descriptive Non-Humean. However, I do not think that this is an appealing option. On the surface, it appears that the denial of (2) requires the denial

of empirical foundationalism. It is, of course, open to the Descriptive Non-Humean to deny empirical foundationalism, but (obviously) I do not find this option appealing. I simply note it here, since the success of this argument in particular has a particularly strong tie to my methodological assumptions. There is a related limitation worth noting. The argument appears to commit us to sense data or qualia, understood in an irreducible way.⁶ I don't find this commitment troubling in the least, but then I have no prior commitments to a rigid metaphysical naturalism or any such view that is hostile to sense data or qualia. Since the assumptions (and commitments) required for this argument are in line with my epistemological and methodological assumptions, for the remainder of this chapter I shall assume that this argument is successful.

The argument I have defended is an epistemological regress argument. If this argument is sound, dispositional monism is false. There have been a number of responses to regress arguments against dispositional monism on behalf of Descriptive Non-Humeans, especially to purely metaphysical regress arguments without commitments to sense data or qualia. In the next subsection, I consider an important objection to such regress arguments. My aim is to show that this important objection is ineffectual against my argument from experience. In the following subsection, I show that considerations from my argument undercut the most important argument against categorical properties—that is, they undercut the most important argument for dispositional monism. These arguments are borrowed from (Bird 2007), as his is a recent, careful, and clear exposition of them.

(Note: The discussion in the following two subsections is somewhat tangential to the central line of argument in this chapter. Some readers may wish to skip to 6.5 and return to this discussion after finishing the rest of the chapter.)

⁶ This same limitation exists for inverted spectrum arguments against dispositional monism. The possibility of an inverted spectrum shows that some properties—namely, color properties—are categorical, since color properties are not defined solely in terms of their relations to other properties (including their causal/nomological relations), but instead stand in brute identity (or, in this case, non-identity) relations to one another. I believe that the empirical foundationalist is in a good position to defend such an argument, but I won't discuss it here because it is similar enough in crucial respects to the argument defended in this section.

6.4.2 The Regress Argument Against Dispositional Monism

I shall now quickly explain a standard metaphysical regress argument against dispositional monism. I do so for the purposes of showing that the argument presented above (in 6.4.1) is immune to the criticisms of the standard regress argument. As its name suggests, the regress argument against dispositional monism attempts to show that dispositional monism leads to a vicious regress. But what sort of regress? As it turns out, there are a number of arguments bearing this title (Bird (2007, Chapters 4 and 6) provides a nice summary). I'm going to focus on one particularly influential version. It is this sort of argument that is responsible for Ellis's (2002) and Fales's (1990, 219–220) rejection of dispositional monism. Before explaining Bird's version of the argument, I'll borrow a few quotes from earlier works that explain the intuitive idea behind it.

Keith Campbell (1976, 93):

Is it possible for anything to be constituted by nothing but causal powers? Whatever the answer is to that question, I doubt very much whether it is possible for *everything* to be constituted by nothing but causal powers. But that seems to be the situation [for dispositional monism]. When one point moves another, all that has been shifted is a power to shift powers to shift... But powers to shift *what*?

Brian Ellis (2002, 171):

If all of the properties and relations that are supposed to be real are causal powers, then their effects can only be characterized by their causal powers, and so on. So causal powers are never manifested. They just produce other causal powers in endless sequence.

(Note: here I would prefer not to say that causal powers are never manifested, but rather that causal powers are never manifested in a way that distinguishes one causal chain from another.)

John Foster (1982, 68):

... there seem to be no physical items in terms of whose behaviour the content of the powers could be specified, and consequently, it seems that, in the last analysis, there is nothing which the powers are powers to do.

I particularly like Bird's own explanation of the argument, and it is this version of the argument (and Bird's responses to it) that I shall engage directly. Note that Bird uses the term

'potencies' to refer to dispositional properties.

Consider the analogy with the circularity of definitions. One way of understanding that problem is that we think of a definition giving meaning to a definiendum because the definiens already has meaning. Definition is a way, as it were, of *passing on* meaning, not of creating it *ab initio*. A system where all meanings were given by definitions would not have any semantic value in the system to be passed around by those definitions. What is needed is some other way of providing meaning, e.g. an ostensive definition linking a word to a thing, so that the system as a whole has some real meaning (semantic value) that can be passed on to the definienda. Similarly one might think that a dispositional statement can only represent some real state of affairs because its stimulus and manifestation conditions are at least possible states of affairs. But if the stimulus and manifestation conditions are themselves mere potencies, there seems to be insufficient reality in the system as a whole. What is needed is that at least some properties are non-dispositional in order to inject some genuine being into the system. As Armstrong goes on to say, 'Perhaps accepting that the purely spatiotemporally properties are categorical will give enough categorical basis to blunt the force of this criticism.' (Bird 2007, 101–2)

Bird offers a two-part response to this argument. First, he tries to show where the argument breaks down.

The clear and simple response for the potency theorist is to deny that dispositions are respectable only if they are reducible to counterfactuals (or something else). The dispositional monist's thesis is that real properties just are dispositions, and are necessarily those dispositions. Thus she will reject immediately the assumption of the objection that dispositions are real only if they are really something else. Once we reject that assumption, the objection falls apart. (Bird 2007, 102)

I am somewhat sympathetic to this response. I say this because I am comfortable (enough) with the idea of a set of languages that consist solely of semantic rules and are completely lacking in any primitive semantic values. Of course, it would be difficult to distinguish one such language from another, and it is not clear what the role of such languages could be. (Bird's idea, I think, will be that each language forms a sort of network or structure, and that the languages can be distinguished by these structures.) But we can set these worries aside. Similarly, one might think that the concept of a world without any categorical properties is perfectly intelligible, though given a pair of such worlds one might not be able to distinguish one from the other. So let's assume

for the sake of argument that Bird's response to the version of the regress argument presented above is legitimate.

That said, I think that the analogy with circularity of definitions is nicely accommodated by my approach to the regress objection, explained in 6.4.1 above. Categorical properties are well-suited to play the role of semantic primitives; dispositional properties are not. Of course, this is how the authors above would interpret things, but I believe that the analogy is stronger than Bird would admit. The language introduced in the context of the interpretation of probability treats Humean predicates (members of Φ) as basic semantic values, some of which are *given in experience* (and the other, more complex ones will be analyzed in terms of those that are given in experience). In contrast, it treats Non-Humean predicates (members of Φ , bare dispositions) and members of Ω (governing laws) as "rules" that restrict the way in which the Humean members of Φ are to be put together with individuals to form consistent sentences (state descriptions). On this framework, only Humean predicates are of the sort that can be given in experience. They are primitives, of course, but this primitiveness is not problematic since it is given in experience. In contrast, it seems that no addition of Non-Humean predicates or of governing laws (semantic rules) will ever give us an analysis of Humean predicates. Sure, different characterizations of Ω and of the Non-Humean members of Φ result in different state descriptions (different possible worlds), but these differences do not suffice to explain differences in possible *experiences* we could have. Thus the Bird's response to the original regress argument explained above does not constitute an adequate response to my argument.

Bird gives a second response. He tries to show that if dispositions are unsuitable candidates for basic semantic values then so are categorical properties. Here is how Bird sets up his response:

To be precise, we may list the essential features that can be attributed to a categorical property thus:

- (a) it is distinct from (i.e. not identical with) other properties;
- (b) it is a universal and thus can have instances;
- (c) for some n it is an n -adic universal. (Bird 2007, 103)

Bird then argues that the dispositional monist's understanding of dispositions satisfies (a), (b), and (c). I don't have any objections to this argument. He continues:

Thus everything attributable to the being of a categorical property is also attributable to the being of a potency. What distinguishes potencies is the *additional* claim that they have (essentially) a dispositional character. Thus there is *more* to the being of an essentially dispositional property than there is to that of a categorical property. In which case the claim that essentially dispositional properties are lacking in reality unless reducible to or explicable in terms of a 'decently real categorical property' is in error. If anything the boot is on the other foot. The thinness of the nature of a categorical property should raise questions about its sufficiency for reality. (Bird 2007, 103)

The problem with Bird's argument is that (a), (b), and (c) do *not* exhaust the essential features of a given categorical property (or, at least, of some categorical properties); together, they don't explain the feature of categorical properties according to which the instantiations of a given categorical property must resemble one another, and the respects in which they must differ from (that is, not resemble) the instantiations of a distinct categorical property. To put things crudely, a given categorical property will have a certain quality in virtue of which its instantiations resemble one another (and it is this quality in virtue of which (a) is true). This quality is brute. Bird might call it a quiddity, and object to this characterization on those grounds (see 6.4.3 below). However, it is precisely that quality of a categorical property—the resemblance relations it entails among its instantiations—that makes it the categorical property it is. According to empirical foundationalism, this is the feature of a categorical property with which we have acquaintance. Conditions (a), (b), and (c) are abstractions from or postulates about this quality of such properties, so contrary to Bird they do not exhaust the being of a categorical property. This is why it makes sense to use the analogy with semantics above; it is the possession of such basic qualities that makes categorical properties suitable candidates for the basic semantic values of predicates in a language. In contrast, since dispositions are intrinsically modal (and thus, from the standpoint of empirical foundationalism, intrinsically unobservable) and since their modal features (plus (a), (b), and (c)) exhaust their being, they are unsuitable as basic semantic values. Bare dispositions lack that which makes categorical properties suitable to play the role of basic semantic values. The latter

have a sort of bruteness about them, but it is the sort that can be given in experience. The former do not, nor can they be given in experience. Thus Bird's definition of categorical properties is incomplete; it is not the case that everything attributable to the being of a categorical property is also attributable to the being of a dispositional property. Once again, we see that Bird's response to this well-known version of the regress argument is ineffectual against the regress argument I defended above.

Summing up, the essential idea is that some semantic values *must* be primitive. How else would we get meaning in a foundationalist framework? Bird treats semantic rules as primitive; why can't we treat values in the same way, especially when we have access to such primitives through immediate experience? It looks like we need both semantic values and semantic rules to generate a language that describes our experience of the world. The upshot, at this point, is that we have a good reason to be skeptical of dispositional monism, and, further, good reason to believe that observed regularities constitute regularities among categorical properties. In 6.4.3 I shall examine an argument in favor of dispositional monism, and I shall argue that the considerations introduced by my argument allow us to easily avoid Bird's criticisms of categorical properties.

6.4.3 Does Dualism Imply Quidditism? Is That Bad?

I shall now discuss Bird's attempt to show that the idea of a categorical property is problematic. Since dispositional monism is the only view according to which there are no categorical properties, this is an argument in favor of dispositional monism.

Before explaining Bird's argument for Dispositional Monism, we need to take another close look at how Bird understands categorical properties. Bird explains,

What we mean by 'categorical' must be understood in negative terms. That is, a categorical property does not confer of necessity any power or disposition. Its existence does not, essentially, require it to manifest itself in any distinctive fashion in response to an appropriate stimulus (Armstrong 1977: 80–3). The categorical versus essentially dispositional distinction is a modal one. To say that a property is categorical is to deny that it is necessarily dispositional (Armstrong *et al.* 1996: 16–17). More generally fundamental categorical properties have no necessary connections with other entities. (Bird 2007, 66–67)

Further,

Indeed, there is very little on [the categoricalist view] to the nature of a given property and certainly nothing that would distinguish it from some other property. The identity and distinctness of properties is a brute fact, not grounded in qualitative differences. (Bird 2007, 3)

I agree mostly with Bird. In denying that a property Q is dispositional we have *not* denied that instantiations of Q must resemble one another in some respect. Recall that, by definition, multiple instantiations of the same property must resemble in some respect. But in what kind of respect could they resemble one another? If it's not modal and not relational (and it isn't), it must be *brute*—that is, primitive. But some primitives are acceptable, especially those which are ultimately grounded in our experiences. Accordingly, my suggestion is that categorical properties are acceptable primitives: either they are given immediately in experience, or they are theoretical postulates whose analysis is ultimately grounded in our immediate experience in the right sort of way (namely, such that they don't constitute bare dispositions).

Thus there is one important respect in which I disagree with Bird. Bird says that there is nothing that would distinguish one categorical property from another. This is simply false. Distinct categorical properties have a distinct “bruteness” to them—the feature in virtue of which instantiations of one property resemble one another but do not resemble the instantiations of a distinct property. There aren't any relational properties that seem to distinguish red and green, but there certainly is something that distinguishes them, and that something is given in immediate experience. So I disagree with Bird when he says that “Categorical properties are all essentially alike—differing only in their mutual distinctness” (2007, 72). This difference of opinion has a great impact on how Bird's arguments are evaluated.⁷

Now that we are aware of this difference of opinion, we can discuss Bird's argument against categorical properties. Bird argues that categorical properties imply the position known as *quidditism*. Quidditism is an analogue of *haecceitism*, which Bird understands as follows: the core of

⁷ Bird says that his view is more or less the view of both David Armstrong (see above) and David Lewis (1986b, 205). However, insofar as Bird and I disagree, I'm not sure whether Armstrong and Lewis are in agreement with me or with Bird.

haecceitism is that the transworld identity of particulars does not supervene on their qualitative features in the given worlds, where qualitative features are understood to be those which exclude properties or relations involving identity (Bird 2007, 71). Accordingly, quidditism is the view that the transworld identity of *properties* does not supervene on the second-order properties they have nor on the qualitative relations they stand in to other properties. The relations here include, among other things, causal and nomological relations. So if quidditism is true, properties are not defined by the causal and nomological networks into which they enter.⁸

So far as I can tell, my view of categorical properties satisfies the definition of quidditism just provided. On my view, the identity of a given property is a sort of brute fact, but one with which we can become acquainted through observing the resemblance relations holding between its instantiations. Since it is brute, the identity of a property doesn't supervene on its second-order properties plus its qualitative relations to other properties (including causal and nomological relations). However, we need to be very careful. Bird claims to provide an argument against quidditism, but his is an argument against two (perhaps closely related) statements, not against quidditism as we have just defined it. Thus, in order to see why Bird's argument against quidditism—and, thus, against the view that some natural properties are categorical—fails, we shall need to carefully distinguish quidditism as it is defined above from quidditism as it is defined by Bird.

Bird provides two definitions of quidditism and argues against both. However, we don't need to look at these arguments carefully to see that my version of quidditism is immune to them.

(QA1) For all fundamental universals F and powers X there is a world where F lacks X . (2007, 71)

(QB1) Two distinct worlds, w_3 and w_4 , may be alike in all respects except that: (i) at w_3 , universal F has powers $\{C_1, C_2, \dots\}$; (ii) at w_4 , universal G has powers $\{C_1, C_2, \dots\}$; (iii) $F \neq G$. (2007,

⁸ We could probably do without the qualifier 'transworld'. If, instead, we defined haecceitism as nothing more than the rejection of the identity of indiscernibles (so that no reference to non-actual worlds is required), I believe that the arguments would apply in more or less the same manner.

72)⁹

Against (QA1), following Black (2000), Bird applies an analogue of Chisholm's (1967) well-known argument against haecceitism. The basic idea is that the (QA1) leads to an absurd view of properties, according to which two worlds can be distinct simply on the grounds that they possess distinct properties, though in those worlds the distinct properties "look" exactly the same, play exactly the same causal/nomological roles, and so on. Against (QB1), Bird argues that it would require us to accept that theories can be multiply realized, and thus that a given theoretical term taken to describe a fundamental property in nature may not refer to a single property, but rather to a group of properties that play the same functional role. In this case, a statement like 'the unique particle with negative unit charge' just won't refer; Bird thinks that that leads to an undesirable skepticism than dispositional monism avoids. Thus he claims, "We do not want our metaphysics of properties to condemn us to necessary ignorance of them. And so we should reject quidditism" (Bird (2007, 78).

My version is immune to both of Bird's arguments because it entails neither (QA1) nor (QB1). First, note that both (QA1) and (QB1) are false if there can be Governing Non-Humean laws of nature. Suppose that there is a synthetic necessary connection (a governing law, a member of Ω) that holds between fundamental categorical properties F and G in all worlds, such that it is necessary that all F s are G s (but that it is not necessary that all G s are F s). Then there cannot be a world in which some F s are not G s, despite the fact that F and G are categorical properties. It follows that, necessarily, objects with F are disposed to be G in all possible conditions. By D, F is therefore a disposition. But F is not a *bare* disposition, since by hypothesis F is a categorical property.¹⁰ Thus, contrary to (QA1), F is a categorical property even though there is no world in which F lacks the relevant power (to become G in all conditions). Thus (QA1) is false. Contrary to (QB1), G lacks a power corresponding to the power that F has—namely, the power to bring about

⁹ The following commentary may be helpful. (QB1) is basically saying the following: exactly the same propositions are true in these worlds, except that any true proposition in one world referring to universal F will be replaced in the other world by a proposition referring to universal G , and vice versa.

¹⁰ I thank Michael Tooley for suggesting that I use this sort of example.

the other property in all conditions; by hypothesis, the possession of *G* does not necessitate the possession of *F*. Thus (QB1) is false.

But for the same reason that (QA1) and (QB1) are false, my version of quidditism—according to which the identity of a categorical property is a brute fact—doesn't entail either (QA1) or (QB1). In the example just given, *F* is a categorical property. Thus *F*'s identity conditions are brute. Thus *F* should be compossible with any given power. But our example has shown that it isn't. If governing laws are possible, categorical properties are not compossible with any given power. Hence the acceptance of categorical properties—properties whose identity is a brute fact—doesn't require the acceptance of (QA1). Similarly, if governing laws are possible, there is no guarantee that the powers of distinct categorical properties are interchangeable—that will depend on the content of the relevant governing laws, and the example above showed that some laws are such that the powers of a categorical property are *not* interchangeable. Thus the acceptance of categorical properties doesn't require the acceptance of (QB1).

The upshot? Bird's arguments against quidditism—and thus against categorical properties—fail, because he doesn't account for the possibility of Governing Non-Humean laws of nature. Once we accept the possibility of such laws, however, we see that the acceptance of categorical properties (and quidditism) entails neither (QA1) nor (QB1). Thus the fact that these principles are false isn't an objection to categorical properties.

Summing up, I've argued that empirical foundationalists have good reason to reject dispositional monism. There is, however, a fair amount of unexplored territory surrounding these arguments. The fairly quick discussion given here does not do full justice to the debate over the nature of natural properties. That being said, I believe that my arguments do place a heavy burden on the dispositional monist; they show that much work is required to make dispositional monism a tenable view (at least from the standpoint of empirical foundationalism). Thus, for the remainder of this chapter I shall assume that dispositional monism is false and that observed regularities constitute regularities among categorical properties.

6.5 The Explanatory Weakness of Descriptive Non-Humeanism

I shall now return to the central question of this essay: can Descriptive Non-Humeanism offer a good explanation of observed regularities? I shall argue that it cannot; it merely relocates the need for explanation by postulating regularities elsewhere.

In the explanation of regularity provided in Section 6.3, the basic idea was that it is necessary that all objects that possess *D* and are in conditions *C* attain *M*. This entails the regularity that all *D* are *M* in *C*. This is a perfectly acceptable explanation of the regularity in question. So what's the problem?

The problem is that *D* is a bare disposition. We don't observe instances of *D*; *D* is an unobservable theoretical property. Though we have a regularity (namely, that *D* are *M* in *C*), we don't have an *observed* regularity. We don't ever observe regularities that can be explained by the Descriptive Non-Humean appeal to bare dispositions, since one of the properties involved in such regularities must always be an unobservable bare disposition. What we need is an explanation of regularities that hold between categorical properties—the properties given in experience—but that sort of explanation has not been provided by this account. In sum, the explanation of regularities provided in 6.3 cannot be extended to explain observed regularities.

An example: Suppose that *water-solubility* is a bare disposition such that, necessarily, anything with it dissolves when placed in water. We know what it's like to observe that something is placed in water. We know what it's like to observe something dissolve. But what is it like to observe the water-solubility of an object? It's not like anything, so far as I can tell; the possession of that property is unobservable. Thus the regularity that everything water-soluble dissolves when placed in water is unobservable. What we want in this case is an observable regularity: for example, that all salt dissolves when placed in water. An explanation for that sort of regularity has not been provided.

I shall now consider a few ways in which the Descriptive Non-Humean might attempt to extend their theory to cover observed regularities. I'm not sure that these options exhaust the

possibilities, but they are the best options I can think of. Suppose that the Descriptive Non-Humean makes the following move: we can explain observed regularities by postulating that there is a categorical property Q that is always co-instantiated by D . (To help keep variables straight, the reader can associate Q with *Qualitative* properties.) This is the natural move to make. We have to get the relevant categorical properties into the picture somehow! Here, the regularity that all D are M in C will also be the regularity that all Q are M in C . The latter regularity is an *observable* regularity. However, there is an obvious problem with this proposal. It gives rise to two options, neither of which is satisfactory: either it is contingent that Q and D are co-instantiated—that is, contingent in the sense that there is no necessary connection binding D to Q —or it is necessary.

If it is contingent that all Q are D then we have just pushed the regularity in need of explanation up one level. The regularity that all Q are D is just as much in need of explanation as the regularity that all Q are M in C . Here, the regularity that all Q are D is doing the crucial work—without that regularity the regularity that all Q are M in C goes unexplained—but the regularity that all Q are D has not been explained. Accepting that all Q are D is contingent—the result of chance, as opposed to the result of some necessary connection between Q and D —is just as damning as accepting that the regularity that all Q are M in C is contingent.

On the other hand, if it is necessary that D and Q are always co-instantiated then one must give an account of the necessity. There seem to be three options, three different accounts of the necessity.

First, one can treat this necessary truth as primitive. That won't work, because the holding of a primitive necessary relation between Q and D constitutes a state of affairs distinct from the regularity that all Q are D ; unlike the state of affairs consisting of a particular's possession of D and Q , the state of affairs consisting of an unanalyzable necessary relation between Q and D isn't a part of the regularity that all Q are D . This, of course, is equivalent to accepting the existence of a governing law relating Q and D , and is thus inconsistent with Descriptive Non-Humeanism.

Second, one can attempt to analyze the necessary truth that all Q are D by postulating a further bare disposition D^* relating Q to D in the following way: necessarily, anything with D^*

has both Q and D . However, in this case we would need to explain the regularity that all Q are D^* —it's not enough to show that all D^* are Q —and the same problem would arise, leading to a vicious explanatory regress. To clarify, the original regularity in need of explanation is that all Q are M in C . We then postulated a disposition, D , such that the possession of D guarantees the possession of M in C . The proposal under consideration asks us to postulate a further disposition, D^* , such that the possession of D^* guarantees the possession of both D and Q . This gets things backwards. Though it explains the regularity that all D^* are D and Q , it does not explain the required regularity—namely, that all Q are D .¹¹

Third, one can attempt to analyse the necessary truth that all Q are D by claiming that D is part of the very concept of Q (suppose for the moment that D and Q are predicates rather than properties in the world). The problem with this strategy is that it makes Q unobservable. On this proposal, that which we observe is merely a property Q^* that is *categorically equivalent* to (that is, looks the same to us as) Q , but does not necessarily involve D . So we can have two objects resemble in respect Q^* (the purely categorical respect) without resembling in respect Q (that is, without resembling in multiple respects Q and D). Thus, since we cannot observe whether an object is Q , this account cannot explain the observed regularities in question; none of those regularities involve Q . (This may sound confusing, but the problem here will become clearer in the context of the example presented below.) Another way to state the problem is to point out that, on this proposal, Q cannot be a categorical property.

In sum, for the account to explain the regularity that everything with Q has M in C we require the extra step that everything with Q has D . The Descriptive Non-Humean cannot treat the connection between Q and D as contingent; the regularity that all Q are D is just as much in need of explanation as the regularity it is supposed to explain. The Descriptive Non-Humean cannot treat the connection between Q and D as a primitive necessary connection; that would be to endorse governing laws. The Descriptive Non-Humean cannot postulate additional bare dispositions to explain the necessary connection between Q and D ; that would lead to a vicious

¹¹ That is, the regularity $\forall x(D^*x \supset (Qx \& Dx))$ does not entail $\forall x(Qx \supset Dx)$.

explanatory regress, since at each higher level there is a new regularity in need of explanation. The Descriptive Non-Humean cannot bind *Q* and *D* together by definition; that makes the possession of *Q* unobservable, and so fails to explain any observed regularity.

An example: Consider the regularity that all salt dissolves in water. According to Descriptive Non-Humeanism, we explain this regularity as follows. There is a property of water-solubility such that, necessarily, everything which possesses water-solubility dissolves when placed in water. All salt possesses water-solubility. That is why we observe the regularity that all salt dissolves when placed in water.

In this explanation, there are four crucial properties. There are three categorical properties identifiable by observation: *salt*, *dissolving*, and *placed in water*. In addition, there is an unobservable dispositional property postulated to explain the regularity that all salt dissolves when placed in water: *water-solubility*. (I am treating *water-solubility* as a bare disposition here, though of course no one thinks that water-solubility is really a bare disposition.) Water-solubility is not directly observable; unlike the other three properties in question, there is nothing that it is like to see it, taste it, touch it, etc. The account in question explains perfectly well why everything with water-solubility dissolves when placed in water, but that's not the observed regularity we set out to explain in the first place. It does *not* explain why all *salt* dissolves when placed in water unless we stipulate that the instantiation of salt always brings with it the instantiation of water-solubility. However, the regularity that all salt has water-solubility is just as much in need of explanation as the regularity that all salt dissolves when placed in water.

We cannot explain the regularity by positing a primitive necessary connection between *salt* and *water-solubility*, because that would be to endorse a governing law and thereby deny Descriptive Non-Humeanism.

We cannot explain the regularity by postulating a new dispositional property *D** that all salt has in virtue of which all salt has water-solubility, since then we will need to explain the regularity that all salt has *D**, and so on ad infinitum. And of course it won't do to say that there is a property, *D**, such that everything with it is both salt and water-soluble, since, once again,

that won't explain the regularity that all salt has water-solubility.

We cannot explain the regularity by stipulating that the concept of *salt* includes the concept of *water-solubility*—that is, by binding the two together by definition such that nothing that fails to possess water-solubility is salt. This strategy fails to account for the nature of our observable experiences; it entails that salt is not a categorical property. We can easily introduce a new concept: *schmalt*, a property which is categorically equivalent to salt but lacking the definitional connection to water-solubility. With this concept on the table, it is easy to see that we cannot actually observe that a given substance is salt; instead, that quality which we observe of a given substance is merely its schmaltiness. We may have thought that we were observing the regularity that all *salt* dissolves in water, but all along we were just observing the regularity that all *schmalt* dissolves in water.

To sum up, although Descriptive Non-Humeanism provides a perfectly good explanation of certain natural regularities (such as the regularity that everything with water-solubility dissolves when placed in water), it cannot provide an explanation of *observed* natural regularities, since it does not permit an explanation of natural regularities between categorical properties, the properties we observe when we observe regularities. There are two natural responses to this problem. First, Descriptive Non-Humeans can reject Dualism in favor of Dispositional Monism and attempt to give an account of how Dispositional Monism accounts for the origin of categorical concepts and for our experience of basic properties as categorical. Second, the dualist can reject Descriptive Non-Humeanism in favor of Governing Non-Humeanism. As I mentioned earlier, I think that there are reasons to prefer the second option to the first, though I haven't provided the arguments for this view in this chapter.

6.6 Conclusion: The Explanatory Power of Governing Non-Humeanism

Governing Non-Humeanism isn't susceptible to the difficulty that plagues Descriptive Non-Humeanism. Though one might reasonably worry that the postulate of governing laws is analogous to the postulate of bare dispositions—in both cases, we are postulating an unobservable state to explain an observable one—there is an important difference between the two.

Governing laws are postulated to be unobservable relations which hold between tuples of properties, but nothing prevents these tuples from consisting of categorical properties only. According to Governing Laws, it is necessary (or probable) that, for any properties F and G , if a nomic relation holds between them then a given regularity obtains between instances of F and G . However, unlike the postulate of bare dispositions, *all* of the relata of such nomic relations may be categorical properties. Thus suitably postulated governing laws entail observable regularities. It is impossible for the Descriptive Non-Humean to give this sort of explanation, since at least one of the relata must be a bare disposition, the possession of which by any object constitutes an unobservable state.¹² In short, the world is structured differently if there are governing laws as opposed to descriptive laws; I have argued here that these structural differences make a big difference in a theory's ability to explain observed regularities.

Suppose we accept governing laws in order to explain the regularity that all Q are D . If we go this route, it seems that we no longer have any need to treat D as a bare disposition; we can reduce D by positing a governing law that holds between Q , M , and C (in which case all Q will possess D , but D will not be a bare disposition; it will be analysed in terms of the laws). Thus it appears that it is (ordinary) bare dispositions, not governing laws, which are dispensable.

To close, I would like to remind the reader that to some extent Governing Non-Humeanism is a compromise position between the popular theories which accept Descriptive Laws and Governing Laws. I agree with Bird, Mumford, and other Descriptive Non-Humeans that some necessity must be taken as basic—namely, the necessary connection between law and regularity. Interestingly, this appears to commit me to the view that fundamental nomic relations are essentially dispositional (since such relations can be shown to satisfy Disposition or something very much like it with relatively little effort). As should be familiar, I have no problem with bare dispositions. The arguments of this chapter are not arguments against bare dispositions. They are simply arguments that bare dispositions cannot do the work of governing laws unless the states of affairs involving bare dispositions constitute governing laws themselves. Thus, as I have argued in

¹² I defend the explanatory power of Governing Non-Humeanism with much greater care elsewhere.

this chapter, I agree with Dretske, Tooley, and Armstrong that governing laws are indispensable. Thus concludes my argument for Governing Non-Humeanism.

Chapter 7

Conclusion

7.1 Final Remarks on Laws of Nature

In Chapter 1, I explained my preferred way of dividing up the logical space concerning theories of laws of nature. In Chapter 2 I set out and defended an empirical method of practicing genuine metaphysics. Chapters 3 through 6 apply this method to the metaphysics of laws of nature. In those chapters, I defended the following theses:

- (1) Empirical foundationalists can give an intelligible account of the concept of synthetic necessity, and therefore can give an intelligible account of the distinction between Humeanism and Non-Humeanism. (Chapter 3)
- (2) Governing Non-Humeanism is a priori at least as likely as Descriptive Humeanism. (Chapter 3)
- (3) Governing Humeanism appears to be internally inconsistent; if it is not, its a priori probability is lower than that of Governing Non-Humeanism, and it does not provide a better explanation of natural regularity than Governing Non-Humeanism. (Chapter 4)
- (4) Governing Non-Humeanism offers a much better explanation of the fact that our world is highly regular than does Descriptive Humeanism. (Chapter 5)
- (5) Governing Non-Humeanism is a priori just as likely as Descriptive Non-Humeanism. (Chapter 6)

(6) Governing Non-Humeanism offers a much better explanation of the fact that the observable world is highly regular than does the most plausible version of Descriptive Non-Humeanism. (Chapter 6)

(1) through (6) entail that the probability of Governing Non-Humeanism in light of observed regularities is greater than that of any of its competitors. The end result is that we have excellent reason to accept the existence of governing laws of nature, where the connection between governing law and regularity is understood to be a synthetic necessary connection. Ultimately, I believe that this reason is decisive. There are no decisive a priori arguments against Governing Non-Humeanism, so the only way for a competing theory to be redeemed would be to provide a posteriori arguments against it. This does not seem at all promising. It is very hard to see how any sort of empirical evidence other than observed regularities could be relevant to theories as broad as those under consideration here. Until our world ceases to be regular, we should accept Governing Non-Humeanism.

This conclusion is extremely broad. I have argued that there are governing laws, but I have not attempted to show that some more carefully specified theory of governing laws is correct. For instance, I haven't argued that laws are atomic relations between universals; I haven't argued that laws are the possession of grand bare dispositions by the world as a whole; I haven't argued that laws are preferences in the mind of God. Relatedly, in contrast to the approach taken by many metaphysicians, I have not arrived at the conclusion that there are governing laws or synthetic necessities by generalization from some particular governing law or synthetic necessity (e.g., as in the case of the generalization from the fact that *necessarily, water is H₂O* to the fact that there is some synthetic necessity). Given my methodology, this should come as no surprise. The method deals in theories that are mutually exclusive and jointly exhaustive. One reason for the generality of the conclusion is the manner in which the theories are defined (but more on this in a moment). Another reason is the generality of the data I selected: the mere fact that the world has a high degree of regularity. For example, we wouldn't expect the generalization that the world is regular

to establish that there is a contingent, external relation between universals F and G that entails (or probabilifies) facts about the distribution of particular states of affairs instantiating F and G . For that, we would want at least an observation that all (or most) F s are G s. I mention this because I don't want to leave philosophers with the impression that the method cannot answer narrower questions in ontology. It hasn't answered narrower questions here because (a) I haven't explored how narrower theories deal with the general observation of natural regularities, and (b) I chose to work with a very broad class of observations. In principle, however, there is no reason to think that the method cannot be applied to the comparison of narrower theories.

I mentioned in the abstract that my argument is a new argument for metaphysically interesting a posteriori necessities. A word on the qualifier *metaphysically interesting* is required. I am of the opinion that the classic a posteriori necessities such as that water is necessarily H_2O or that salt is necessarily water soluble are metaphysically *uninteresting*. I don't have the time to defend this claim here (and it isn't required for the claim that my argument establishes the existence of metaphysically interesting a posteriori necessities), though my discussion of the example of salt and water-solubility at the end of Chapter 6 hints at my reasons. The central idea is that such a posteriori necessities have no implications for our possible experiences of the world. How could they? They involve the mere necessity of identity. (I'm presented with a schmalty substance; I'm not allowed to infer that it is water soluble unless I also possess knowledge that all schmalt is salt, but that isn't the a posteriori necessity in question.) It is this sense in which I think such necessities are uninteresting (they may well be philosophically interesting for other reasons). In contrast, I have argued that we have empirical reasons to accept the existence of governing laws of nature. These laws do have important implications for our possible experiences of the world. Thus, this is not merely a new argument for the same old a posteriori necessities of identity; it is an argument for a more powerful species of a posteriori necessity.

7.2 Final Remarks on Method

To conclude, I want to make some general remarks on the method and discuss its implications for issues in contemporary metaontology.

7.2.1 Relation to Quine's Empirical Method

I shall now discuss the relationship between my method and a more common method of practicing metaphysics that is often assumed to be empirically respectable.

The method in question is due to Quine—its essential features may be gleaned from Quine's (1948) and (1951)—and may be described as follows:

Contemporary Quinean Metaphysics (CQM): Metaphysical theorizing, like scientific theorizing, consists of selecting the best set of theoretical beliefs that undergird, support, and fit nicely within our web of pretheoretical¹ background beliefs—that is, our logical beliefs, ordinary beliefs, beliefs concerning immediate experiences, and so on. Potential criteria for selection of the best set include (but may not be limited to) considerations of simplicity, unity, and explanatory strength.²

As stated, Quine's method is not fully specified. In order to apply the method, we need to have some idea of what constitutes the best theory or, more broadly, the best web of beliefs. We need rules that give us some guidance as to how the criteria for theory choice are to be applied. An ideally specified method would include (a) a set of weighted background beliefs and (b) a set of weighted criteria for theory choice (where the weights tell us the importance of each type of belief or criterion).

Unfortunately, it is far from clear how an empiricist is supposed to defend a specified version of CQM method that says anything at all about (a) and (b). If CQM were merely intended to provide pragmatic reasons for theory choice—that is, if it were merely intended to provide a

¹ 'Pretheoretical' can refer to beliefs prior to *metaphysical* theorizing only if we like, so the method allows for the inclusion of non-metaphysical theoretical beliefs such as scientific beliefs in the set of background beliefs.

² See (Sider 2009, 385) for a related statement of CQM.

useful, simple way of systematizing the ontological commitments of our closely-held background beliefs—then it would not be terribly difficult to justify a specified method. For instance, it is easy to show that *simplicity* is a pragmatic virtue of a theory. Indeed, it appears that Quine interprets his own method as merely providing pragmatic reasons for theory choice.³ But we want more than that. We want to know what the world is really like. The goal of metaphysics, as I understand it, is to investigate *reality*, not merely our use of language or our way of forming concepts. Since we want to know which theory is *true*, a method of practicing metaphysics must provide *epistemic* reasons for theory choice.

Unfortunately, providing epistemic justification for a carefully specified version of CQM (that is, for a given assignment of weights to background beliefs and criteria for theory choice) is very difficult. For example, Sider (2009, 385) mentions some questions that shed doubt on CQM:

What justifies the alleged theoretical insights? Are criteria that are commonly used in scientific theory choice (for example, simplicity and theoretical integration) applicable in metaphysics? How can these criteria be articulated clearly? And what hope is there that the criteria will yield a determinate verdict, given the paucity of empirical input?

Questions like these raise serious problems for CQM. Proponents of CQM who claim to provide epistemic reasons for theory choice must be able to answer them.

My point here is not to argue that the nowadays-popular criteria for theory choice cannot be supported (though I do have serious doubts that empiricists can assign any very basic epistemic status to criteria of simplicity, unity, and the like). I certainly do not wish to argue that the general epistemic barriers to CQM are insuperable. I have merely noted that there are epistemic challenges to CQM, and that contemporary metaphysicians ought to be concerned with these challenges. Rather, I highlight these challenges because the method defended in this project constitutes a (partial) response to them. I have justified an empirical method of practicing meta-

³ Quine recommends the view that we have merely *pragmatic* reasons for theory choice in (1951). See also (Price 2009) and (Soames 2009) for arguments that Quine is to be interpreted in this manner. In contrast, (Sider 2009) and (van Inwagen 2009) both interpret Quine's method as providing epistemic reasons for theory choice. This is not to say that they interpret Quine as providing such reasons, only that they take the method to provide such reasons. Manley (2009, 3–4) refers to these realist interpretations as *mainstream metaphysics*—an indicator of the popularity of this position.

physics that fits the general model of CQM (though in a rather different way than Quine would have wanted). The defense of my method is a defense of CQM in the sense that my method assigns weights to background beliefs and to criteria for theory choice. Since the general statement of CQM doesn't tell us how to choose the best web of beliefs, it is consistent with foundationalism and semantic atomism, even though the web metaphor is usually associated with coherentism and semantic holism. The method I have advocated is one way of explaining how the best web of beliefs is to be chosen, and in this sense it is clearly a version of CQM. (It is not surprising that my method can be described in a such a way that it satisfies CQM; CQM is so general that it provides almost no guidance at all.)

However, when we examine my method's way of assigning weight to background beliefs and criteria for theory choice it becomes apparent that my method is quite revisionary. It certainly differs from Quine's approach to metaphysics, even if we pretend for the moment that Quine intended his method to give epistemic reasons for selecting the best web of beliefs. My method holds that certain points of the web are firmly grounded, not subject to revision no matter the theoretical benefits. Analytic propositions and those propositions justified through immediate experience are never to be compromised; these foundational nodes are not open to revision. In addition, my method assigns no weight to certain sorts of background beliefs: beliefs justified by rational intuition, ordinary beliefs, commonsense beliefs, scientific beliefs, and the like. My method assigns no basic weight to certain criteria for theory choice: unity, integration with other domains, simplicity, and the like.⁴ On my method, the metaphysician approaches her trade by discovering analytic linking principles and comparing the antecedents of those principles with her immediate experiences. The analytic principles are taken to be unrevisable parts of her web of beliefs, as are her immediate experiences. The result is that there is no need (and perhaps no room) for basic principles concerning unity, simplicity, and the like.⁵

⁴ Some of these criteria do receive an indirect weight as a result of other components of the method. For instance, a very limited version of Occam's razor may be derived from the version of the logical interpretation of probability I favor.

⁵ Given these differences, my guess is that (Manley 2009, 4) would characterize me as a *reformer* rather than as a *mainstream metaphysician*.

7.2.2 Contemporary Debates in Metaontology

Finally, I want to discuss how this method bears on other contemporary debates in metaontology.

I have focused on answering epistemic challenges to the practice of empirical metaphysics. However, semantic objections to the general practice of metaphysics have become quite popular. The latter object to realist interpretations of CQM on the grounds that metaphysical disputes are either meaningless or trivial.⁶ Note, however, that if one follows the method defended here, such objections lose force for the metaphysical issues receptive to this method. If a metaphysical dispute can be framed in accordance with this method, that dispute is both meaningful and non-trivial, since the method sets out clear conditions for defining the theories such that they are required to differ in either their a priori probabilities or their observational consequences. For example, I have argued that the method can be applied to the metaphysics of laws of nature in such a way that different theories of laws have very different implications for our immediate experiences (namely, that governing laws make it likely that our experiences will be full of regularities, and that non-governing accounts of laws do not). How could theories differing in this way fail to differ in their meaning? The application of my method to the metaphysics of laws of nature shows that there are clear cases for which the semantic objections fail.

There is also a more general implication of this method. It is the breadth of the theories employed by my method that makes genuine metaphysics possible (recall that without Step 4, the constraint on defining theories, the method does not allow for genuine confirmation of theories, and the scope of the method is thus extremely limited). For this reason, the role of defining theories is given a central position within the process of theorizing. Any old statement of a theory won't do, because not just any old theory (or set of theories) may be "plugged in" to this method. The theories being compared must be mutually exclusive and jointly exhaustive, and in almost all cases they must differ with respect to their observable consequences. Thus, the applicability of

⁶ Hirsch (2009) argues that some metaphysical disputes are meaningless. Thomasson (2009) argues that some questions in ontology are either unanswerable or trivial.

this method will depend on our ability to define theories in accordance with these constraints.

Some issues in metaphysics seem to be ripe for such treatment, such as problems in the metaphysics of laws of nature explained in this project. In addition, a number of classical arguments in the area of philosophical theology might be run in accordance with this method, such as teleological arguments and evidential versions of the problem of evil (provided that proper sense can be made of the ethical components of such arguments). On the other hand, other issues appear intractable. The method does leave the door open for a type of epistemic deflationism concerning certain metaphysical issues. It may be the case that, for some metaphysical issues, there is no way of parsing the theories such that (a) there are observations relevant to the dispute, (b) theories differ in their observational consequences, and (c) the theories are such that it is possible for us to determine something about their a priori probabilities. The standard theories in the debates over composition and colocation do not seem amenable to this treatment (at least not to me), since the relevant theories do not appear to differ in their observable consequences. Further, the theoretical criteria of my method seem insufficient to select one theory over another. Thus, the method appears to support Bennett's (2009) argument that we simply aren't in the epistemic position to choose between the standard theories concerning composition and colocation. The debates may be meaningful, but it is difficult to see how my method (or anything like it) could be applied to resolve them. (Note that my method does not suggest that the debates are in fact meaningful, so it is also consistent with a semantic deflationism concerning certain disputes.)

As in the other sections, everything noted in this final section is contingent upon the success of my method. It needs to be shown that empirical foundationalism is consistent with the assumption that we have justification for accepting some system of logic; similarly, it needs to be shown that some method of assigning objective epistemic probabilities is correct; and empirical foundationalism must be established as the correct epistemic theory. However, none of these tasks are strictly required for the demonstration that metaphysics can be informed by our immediate experiences of the world. Even if these assumptions are false, and even if we do possess synthetic a priori reasoning, metaphysics need not be an exercise in synthetic a priori reasoning *only*, nor

an exercise in interpreting the commitments of our best scientific theories *only*. Metaphysics can be informed by our immediate experiences.

Bibliography

- Anscombe, G.E.M. 1971. Causality and Determination. Cambridge: Cambridge University Press.
- Armstrong, David. 1983. What is a Law of Nature? Cambridge: Cambridge University Press.
- . 1988. "Reply to van Fraassen." Australasian Journal of Philosophy 66:224–229.
- . 1999. "Reply to Ellis." In H. Sankey (ed.), Causation and Laws of Nature, 43–48. Dordrecht: Kluwer.
- Beebe, Helen. 2000. "The Non-Governing Conception of Laws of Nature." Philosophy and Phenomenological Research 61:571–594.
- Bennett, Karen. 2009. "Composition, Colocation, and Metaontology." In Ryan Wasserman David Chalmers, David Manley (ed.), Metametaphysics, 38–76. New York: Oxford University Press.
- Bird, Alexander. 2005. "The Ultimate Argument Against Armstrong's Contingent Necessitation View of Laws." Analysis 65:147–155.
- . 2007. Nature's Metaphysics: Dispositions, Laws and Properties. Oxford: Oxford University Press.
- Black, Robert. 2000. "Against Quidditism." Australasian Journal of Philosophy 78:87–104.
- Campbell, Keith. 1976. Metaphysics: An Introduction. Encino, CA: Dickenson.
- Carnap, Rudolf. 1952. The Continuum of Inductive Methods. Chicago: University of Chicago Press.
- . 1959. "The Elimination of Metaphysics Through Logical Analysis of Language." In A.J. Ayer (ed.), Logical Positivism, 60–81. Macmillan.
- . 1962. The Logical Foundations of Probability. Chicago: University of Chicago Press, 2nd edition.
- . 1966. Philosophical Foundations of Physics. New York: Basic Books.
- Carroll, John. 1994. Laws of Nature. Cambridge: Cambridge University Press.
- Carroll, John (ed.). 2004. Readings on Laws of Nature. Pittsburgh: University of Pittsburgh Press.

- Chisholm, Roderick. 1967. "Identity Through Possible Worlds: Some Questions." Nous 1:1–8.
- Dretske, Fred. 1977. "Laws of Nature." Philosophy of Science 44:248–268.
- Earman, John and Roberts, John T. 2005a. "Contact with the Nomic: A Challenge for Deniers of Humean Supervenience about Laws of Nature Part I: Humean Supervenience." Philosophy and Phenomenological Research 71:1–22.
- . 2005b. "Contact with the Nomic: A Challenge for Deniers of Humean Supervenience about Laws of Nature Part II: The Epistemological Argument for Humean Supervenience." Philosophy and Phenomenological Research 71:253–286.
- Ellis, Brian. 2001. Scientific Essentialism. Cambridge: Cambridge University Press.
- . 2002. The Philosophy of Nature. Chesham: Acumen.
- Fales, Evan. 1990. Causation and Universals. New York: Routledge.
- Fine, Kit. 1978. "Model Theory for Modal Logic Part I: The De Re / De Dicto Distinction." Journal of Philosophical Logic 7:125–156.
- Forbes, Graeme. 1985. The Metaphysics of Modality. Oxford: Oxford University Press.
- Foster, John. 1982. The Case for Idealism. London: Routledge and Kegan Paul.
- . 1982–83. "Induction, Explanation, and Natural Necessity." Proceedings of the Aristotelian Society 83:87–101.
- Goodman, Nelson. 1955. Fact, Fiction, and Forecast. Cambridge, Mass.: Harvard University Press.
- Hintikka, Jaakko. 1965. "Towards a Theory of Inductive Generalisation." In Y. Bar Hillel (ed.), Logic, Methodology and Philosophy of Science, 274–288. Amsterdam: North Holland.
- . 1966. "A Two-Dimensional Continuum of Inductive Methods." In Jaakko Hintikka and Patrick Suppes (eds.), Aspects of Inductive Logic, 113–32. Amsterdam: North Holland.
- Hirsch, Eli. 2009. "Ontology and Alternative Languages." In David Chalmers, David Manley, and Ryan Wasserman (eds.), Metametaphysics, 231–259. New York: Oxford University Press.
- Howson, Colin and Urbach, Peter. 1993. Scientific Reasoning: The Bayesian Approach. Chicago: Open Court, 2nd edition.
- . 2006. Scientific Reasoning: The Bayesian Approach. Chicago: Open Court, 3rd edition.
- Kripke, Saul. 1963. "Semantical Considerations on Modal Logic." Acta Philosophica Fennica 16:63–72.
- Lange, Marc. 2009. Laws and Lawmakers. Oxford: Oxford University Press.
- Lewis, David. 1970. "How to Define Theoretical Terms." The Journal of Philosophy 67:427–446.
- . 1973. Counterfactuals. Cambridge, Mass.: Harvard University Press.
- . 1986a. "Against Structural Universals." Australasian Journal of Philosophy 64:25–46.

- . 1986b. On the Plurality of Worlds. Oxford: Blackwell.
- . 1994. "Humean Supervenience Debugged." Mind 412:473–490.
- Loewer, Barry. 1996. "Humean Supervenience." Philosophical Topics 24:101–127.
- Mackie, J.L. 1977. Ethics. London: Penguin Books.
- Manley, David. 2009. "Introduction: A Guided Tour of Metametaphysics." In David Chalmers, David Manley, and Ryan Wasserman (eds.), Metametaphysics, 1–37. New York: Oxford University Press.
- Mellor, D.H. 2005. Probability: A Philosophical Introduction. New York: Routledge.
- MicKitrick, Jennifer. 2003. "The Bare Metaphysical Possibility of Bare Dispositions." Philosophy and Phenomenological Research 66:349–369.
- Mikkelsen, Jeffrey M. 2004. "Dissolving the Wine/Water Paradox." British Journal for the Philosophy of Science 55:137–145.
- Monton, Bradley. 2006. "God, Fine-Tuning, and the Problem of Old Evidence." British Journal for the Philosophy of Science 57:405–424.
- Moore, G.E. 1925. "A Defence of Common Sense." In J.H. Muirhead (ed.), Contemporary British Philosophy (2nd Series), 193–223. London: Allen and Unwin.
- Mumford, Stephen. 2004. Laws in Nature. London: Routledge.
- Plantinga, Alvin and Tooley, Michael. 2008. Knowledge of God. Malden, Mass.: Blackwell.
- Popper, Karl. 1959. The Logic of Discovery. New York: Basic Books.
- Price, Huw. 2009. "Metaphysics after Carnap: The Ghost Who Walks?" In David Chalmers, David Manley, and Ryan Wasserman (eds.), Metametaphysics, 320–346. New York: Oxford University Press.
- Prior, Elizabeth, Pargetter, R., and Jackson, Frank. 1982. "Three Theses About Dispositions." American Philosophical Quarterly 19:251–257.
- Quine, W.V.O. 1948. "On What There Is." Review of Metaphysics 2:21–38.
- . 1951. "Two Dogmas of Empiricism." The Philosophical Review 60:20–43.
- Ramsey, F.P. 1931. "Theories." In R.B. Braithwaite (ed.), The Foundations of Mathematics. London: Routledge and Kegan Paul.
- Russell, Bertrand. 1912. The Problems of Philosophy.
- Schaffer, Jonathan. 2008. "Causation and Laws of Nature." In Theodore Sider, John Hawthorne, and Dean Zimmerman (eds.), Contemporary Debates in Metaphysics, 82–107. Oxford: Blackwell.
- Sider, Theodore. 1992. "Tooley's Solution to the Inference Problem." Philosophical Studies 67:261–275.

- . 2001. Four Dimensionalism: An Ontology of Persistence and Time. Oxford: Oxford University Press.
- . 2009. "Ontological Realism." In David Chalmers, David Manley, and Ryan Wasserman (eds.), Metametaphysics, 384–423. New York: Oxford University Press.
- Soames, Scott. 2009. "Ontology, Analyticity, and Meaning: the Quine–Carnap Dispute." In David Chalmers, David Manley, and Ryan Wasserman (eds.), Metametaphysics, 424–443. New York: Oxford University Press.
- Sober, Elliot. 2003. "The Design Argument." In N. Manson (ed.), God and Design: The Teleological Argument and Modern Science, 27–54. London: Routledge.
- Stalnaker, Robert. 1968. "A Theory of Conditionals." In Nicholas Rescher (ed.), Studies in Logical Theory, volume 2 of Monograph Series, 98–112. American Philosophical Quarterly, Oxford: Blackwell.
- Swinburne, Richard. 1980. "Reply to Shoemaker." In L.J. Cohen and Mary Hesse (eds.), Applications of Inductive Logic, 313–320. Oxford: Oxford University Press.
- Thomasson, Amie. 2009. "Answerable and Unanswerable Questions." In David Chalmers, David Manley, and Ryan Wasserman (eds.), Metametaphysics, 444–471. New York: Oxford University Press.
- Tooley, Michael. 1977. "The Nature of Laws." Canadian Journal of Philosophy 667–698.
- . 1987. Causation: A Realist Approach. Oxford: Oxford University Press.
- . 1999. "Introduction." In Michael Tooley (ed.), Metaphysics, volume 3. Garland.
- . Forthcoming. Causation: Fundamental Issues. Oxford University Press.
- van Fraassen, Bas. 1989. Laws and Symmetry. Oxford: Clarendon Press.
- van Inwagen, Peter. 1998. "Modal Epistemology." Philosophical Studies 92:67–84.
- . 2009. "Being, Existence, and Ontological Commitment." In David Chalmers, David Manley, and Ryan Wasserman (eds.), Metametaphysics, 473–506. New York: Oxford University Press.