

THE BEST PREDICTIVE SYSTEM ACCOUNT OF LAWS OF NATURE

Chris Dorst

A dissertation submitted to the faculty at the University of North Carolina at Chapel Hill in partial fulfillment of the requirements of the degree of Doctor of Philosophy in the Philosophy Department in the College of Arts and Sciences.

Chapel Hill
2018

Approved by:

Marc Lange

John Roberts

Matthew Kotzen

Laurie Paul

Edward Hall

© 2018
Chris Dorst
ALL RIGHTS RESERVED

ABSTRACT

Chris Dorst: The Best Predictive System Account of Laws of Nature
(Under the direction of Marc Lange)

This dissertation develops a novel theory of laws of nature in the “Best System” tradition, with the express aim of making sense of why creatures like us are interested in discovering the laws. My theory draws inspiration from David Lewis’s famous Best System Account of laws. Lewis’s account has two basic elements: the “Humean base” and the “Nomic Formula.” The laws, according to Lewis, are the results of applying the Nomic Formula to the Humean base. My account preserves this overall structure of Lewis’s view, but I disagree with him both about what sorts of facts constitute the Humean base, and about the nature of the Nomic Formula itself. In the first two chapters of my dissertation, I develop objections to Lewis’s explications of these elements, and on the basis of these objections, I propose alternative accounts of the Humean base (Chapter 1) and the Nomic Formula (Chapter 2). In short, my view is that the laws of nature are the principles of the most predictively useful systematization of the totality of macroscopic phenomena. I call the resulting view the “Best Predictive System Account” of laws. In Chapter 3, then, I attempt to explain why the laws tend to be held fixed in counterfactual reasoning. I do so by arguing that, if the Best Predictive System Account is correct, creatures like us would naturally hold fixed the laws in counterfactual reasoning for purposes of figuring out facts about the actual world.

ACKNOWLEDGEMENTS

This dissertation has been improved by conversations and interactions with a number of people, to whom I am very grateful.

Thanks first and foremost to Marc Lange, who has always been incredibly willing to prioritize my requests for advice and comments. Marc's extensive, nearly-immediate feedback at every stage of this process has not only improved this dissertation tremendously; it has also made me a better philosopher, as well as a more thorough, careful thinker. I could not have asked for a better advisor, mentor, and role model.

Thanks also to Ned Hall, who was tremendously accommodating while I spent time working with him at Harvard, and who has been singularly influential in helping me find a philosophical approach that perfectly suits my sympathies. Ned's enthusiasm for philosophy is infectious, and my excitement about my work swells every time I talk with him.

Thanks to John Roberts, first for always finding the weakest parts of my arguments, and then for invariably helping me come up with numerous creative responses. In my conversations with John, it was not uncommon for me to feel both daunted and reassured at the same time—that I was out of my depth, but also that it didn't matter, because he was going to help me get to where I needed to go anyway. My conversations with John are some of the most interesting and absorbing I've ever had.

Thanks to Matt Kotzen for his enthusiasm and insight about this project, and for suggesting numerous connections to other areas of active research. Matt also consistently helped me situate my project within a broader philosophical perspective, and provided invaluable advice

about structuring my papers and crafting compelling narratives. My paper writing improved tremendously under his guidance.

This dissertation has also benefitted from numerous conversations and comments from Kevin Dorst, Stanley Dorst, Laurie Paul, Ram Neta, Rob Smithson, Krasi Filcheva, Kyle Driggers, Alan Nelson, Jerry Postema, Mike Hicks, Erica Shumener, Barry Loewer, David Albert, Christian Loew, Brad Skow, Isaac Wilhelm, Finnur Dellsén, Alfredo Watkins, Carla Merino-Rajme, Alex Worsnip, Zanja Yudell, Heather Demarest, Elizabeth Miller, and Christopher Stephens. Thanks also to Douglas Stalker for much encouragement and motivation.

I would also like to thank participants in the 2017 Pacific APA colloquium session on laws of nature, the 2018 Eastern APA colloquium session on laws of nature, and the 2017 UC Boulder Conference on the history and philosophy of science.

Thanks also to my parents, Stanley Dorst and Margaret Yoest, for supporting me in so many ways throughout this process. And to my dogs, Maple and Mocha, for keeping me company as I worked.

Lastly, thanks to Kayla Dorst for supporting me, for patiently going over my arguments and ideas, for helping me rehearse various presentations, and for pushing me to think about my work from new perspectives. I could not have done this without her.

TABLE OF CONTENTS

LIST OF FIGURES.....	viii
LIST OF ABBREVIATIONS AND SYMBOLS.....	ix
CHAPTER 1: INTRODUCTION.....	1
CHAPTER 2: WHAT IS THE HUMEAN BASE?	4
1. Introduction.....	4
2. Lewis's Best System Account.....	5
2.1 Characterizing the BSA.....	5
2.2 Attractions of the BSA.....	8
3. A Tension in the BSA.....	19
4. Reconceptualizing the Humean Base.....	31
4.1 The Geiger-Marsden Experiment.....	33
4.2 The Discovery of the Omega Minus Baryon.....	37
5. A Sketch of the View.....	48
6. Conclusion.....	57
CHAPTER 3: THE NOMIC FORMULA.....	60
1. Introduction.....	60
2. The LOPP.....	62
3. Problems with Lewis's Formula.....	64
4. A Pragmatic Account of the Nomic Formula.....	76

4.1 Informative Dynamics.....	77
4.2 Wide Applicability.....	79
4.3 Spatial Locality.....	81
4.4 Temporal Locality.....	85
4.5 Spatial, Temporal, and Rotational Symmetries.....	86
4.6 Predictively Useful Properties.....	87
4.7 Simplicity.....	91
5. Why are the Laws So Useful?	95
5.1 Predictive Utility and Nonreductionism.....	95
5.2 Predictive Utility and Reductionism.....	99
5.3 Why the LOPP is Misleading.....	102
6. Conclusion: Why we Care about the Laws.....	104
CHAPTER 4: WHY DO THE LAWS SUPPORT COUNTERFACTUALS?.....	106
1. Introduction.....	106
2. The Altered States Recipe.....	107
3. Counterfactuals as a Guide to Actuality.....	115
4. Humean Ramifications.....	124
5. Conclusion: Does God Think about Counterfactuals?.....	128
CHAPTER 5: CONCLUSION.....	129
APPENDIX.....	138
REFERENCES.....	147

LIST OF FIGURES

Figure 1 – Geiger-Marsden Apparatus.....	34
Figure 2 – The Omega Minus Experiment.....	38

LIST OF ABBREVIATIONS AND SYMBOLS

\Rightarrow	Counterfactual Arrow
BSA	Best System Account
BPSA	Best Predictive System Account
LOPP	Limited Oracular Perfect Physicist
PH	Past Hypothesis
SP	Statistical Postulate

CHAPTER 1: INTRODUCTION

This dissertation aims to develop a novel philosophical account of laws of nature. I call this account the “Best Predictive System Account,” or “BPSA,” for short. As the name suggests, it constitutes a (significant) modification of David Lewis’s Best System Account of laws, or “BSA.” It is, then, a thoroughly “Humean” or “reductionist” account, just as Lewis’s is, but it is not meant to appeal merely to fans of the BSA. In my view, the attractions of reductionism have not yet been fully articulated, and part of my aim here is to strengthen the appeal of a reductionist metaphysics by showing just how strong a theory it can produce. Indeed, I hope that the benefits of the BPSA convince even the most ardent nonreductionists to consider joining the reductionist camp.

Since my view is a modification of Lewis’s BSA, this dissertation is organized around the structure of that view. The BSA itself has two basic elements, which I call the “Humean base” and the “Nomic Formula.” The laws, according to Lewis, are the results of applying the Nomic Formula to the Humean base. The BPSA preserves the basic structure of Lewis’s view, but it differs both in what sorts of facts constitute the Humean base, and in the content of the Nomic Formula itself. Thus, in the first two chapters of this dissertation, I develop objections to Lewis’s explications of these elements, and on the basis of these objections, I propose alternative accounts of the Humean base (Chapter 1) and the Nomic Formula (Chapter 2). I will summarize these discussions in turn.

Briefly put, Lewis thinks of the Humean base as the totality of fundamental facts about the world.¹ The fundamental facts, he thinks, are delivered to us from physics. So one such fact might be the following: “Spacetime point s is occupied by a charm quark with x -spin up.” The laws are then conceived as principles systematizing the totality of these fundamental facts.

There are a number of problems with this conception of the Humean base, many of which have been raised by other philosophers. In Chapter 1 §3, I review some of these problems and also introduce one of my own, which I immodestly suggest is the deepest problem with Lewis’s conception. For reasons that will become clear later, this discussion motivates a departure from the idea that the Humean base consists of the totality of particular *fundamental* facts. Instead, I develop a view whereby the Humean base consists of the totality of particular facts about *macroscopic phenomena*. Part of the motivation for this move is that it more accurately reflects the structure of scientific practice, in which the laws and fundamental kinds are posited *jointly*, in response to observations about the behaviors of macroscopic entities.

The Nomic Formula, according to Lewis, is an operation that gets applied to the Humean base in order to generate the laws. Lewis says that the Nomic Formula involves a balance of two desiderata, simplicity and strength, aimed at achieving an *efficient summary* of the entire Humean base. In Chapter 2, I pose an objection to this characterization of the Nomic Formula: roughly, if it were correct, it would be unclear why creatures like us care about discovering the laws. On the basis of this objection, I develop an improved Nomic Formula that involves a balance of different desiderata which are aimed at achieving a set of principles that is maximally *predictively useful* to creatures like us. The result is a theory of laws that has a number of

¹I review the specifics of Lewis’s theory in much more detail later on. In particular, Chapter 1 §2.1 reviews the structure of the Best System Account in much more detail, and §2.2 discusses some of the main motivations for that structure.

benefits. Most notably, this theory explains why creatures like us are so invested in discovering the laws in the first place, and it also generates laws with the same sorts of formal features that we find in actual scientific practice.

Together, Chapters 1 and 2 constitute the core of the Best Predictive System Account. In short, my view is that the laws of nature are the principles of the most predictively useful systematization of the totality of the macroscopic phenomena.

In Chapter 3, I use this conception of laws to explain why the laws tend to be held fixed in counterfactual reasoning. Of course, it is widely agreed *that* the laws of nature exhibit a peculiar degree of counterfactual resilience, and many accounts of laws and counterfactuals have been designed to accommodate this fact. But here I aim to understand *why* the laws are counterfactually resilient in this manner. A satisfying explanation of the counterfactual resilience of the laws requires a rationale for our engagement in counterfactual reasoning. Here I argue that creatures like us engage in counterfactual reasoning to figure out facts about the actual world. In particular, when we are presented with certain forms of evidence that we did not anticipate, we use counterfactual reasoning to figure out how this evidence bears on certain hypotheses of interest. Roughly, we ask: “Would we have obtained this evidence if such-and-such a hypothesis were true?” We then construct a hypothetical scenario in which the hypothesis in question *is* true, and see whether it leads, via the laws of nature, to the observed evidence. If it does, then the observed evidence favors that hypothesis; if it does not, then the observed evidence counts against that hypothesis. It is crucial here that the laws of nature are held fixed in this form of reasoning, because they allow us to generate a ‘pseudo-prediction’ that tests whether the hypothetical scenario leads to the observed evidence. Thus, given our aims in counterfactual

reasoning, it is *to be expected* that creatures like us would hold fixed our predictive principles (read: “the laws of nature”) when we engage in it.

My view thus situates the laws—and by extension, other nomological concepts such as counterfactuals—in a thoroughly *pragmatic* context. Thus I expect that many people will think the account is too anthropocentric. They will say that any account that makes the laws derivative upon our practical concerns just *can’t* be right, for the laws are purely objective features of reality that don’t care about us and our interests.

I flag this objection now only to show that I am well aware of it. Indeed, an earlier version of myself would have criticized the Best Predictive System Account on exactly these grounds. But I now think that this objection amounts to little more than table-pounding, and that there are powerful arguments motivating the sort of pragmatic approach that I am taking here. I will raise these arguments when it becomes appropriate throughout the ensuing chapters, especially in Chapter 2.² For now, though, I will merely have to plead the reader’s patience.

² See especially Chapter 2, §3, §5, and §6.

CHAPTER 2: WHAT IS THE HUMEAN BASE?

1. Introduction

Perhaps the most famous philosophical account of laws of nature is David Lewis's Best System Account. Anyone who does serious philosophical work on laws has to be familiar with it, if not in detail then at least in abstract. Partly for this reason, it will be convenient to introduce my own account using Lewis's as both a guide and a foil.

Now it turns out—and this will be one of my main points in this chapter—that there is a serious tension hidden within Lewis's account. It arises because of two incongruent ways in which Lewis attempts to hold his theory accountable to actual scientific practice. This tension, however, has gone largely unnoticed in the philosophical literature on laws. I suspect this is because it only becomes apparent when we articulate Lewis's theory in a particular way, i.e. when we emphasize certain aspects of his view that are not normally emphasized. One of my goals in this chapter is to make this tension apparent, and show why it is problematic.

However, I will not be content to just point this out, for I am actually very sympathetic to the overall Humean picture.³ So my other goal in this chapter is to fix that picture—to resolve the tension. To do so, I think we need to reconceptualize the “Humean base” that figures into Lewis's account. In brief, my proposal will be that we need to think of the Humean base as the totality of facts that scientific theorizing about the laws of nature is supposed to systematize and explain.

³There is a real question of whether the view I end up advocating counts as “Humean.” I think it does, but I shall delay discussion of this comparatively unimportant issue until the conclusion of this chapter.

But I am jumping the gun. First I need to explain Lewis's Best System Account of laws and why it's attractive (§2). Then I need to elucidate the tension in his account, and show why it deserves our attention (§3). Only after both of these have been accomplished can I appropriately motivate my positive proposal for the Humean base (§4) and sketch the overall picture provided by my view (§5). I conclude (§6) by defending my branding of this view as "Humean," and an Appendix addresses some outstanding objections.

2. Lewis's Best System Account

2.1 Characterizing the BSA

The canonical formulation of Lewis's Best System Account (BSA) goes roughly as follows.⁴ Fundamentally, all that there is to the world is a vast mosaic of particular matters of fact, which Lewis calls the "Humean base" (or variously, the "Humean mosaic"). The constituents of these facts are discovered by fundamental physics, which provides us with a catalog of the fundamental entities and the (perfectly) natural properties. Thus according to current physics, the fundamental entities could be some sort of particles, like quarks and leptons, and the perfectly natural properties could be things like *x*-spin, *y*-spin, mass, charge, and so on. For brevity, I will refer to all of these—the kinds of fundamental entities and perfectly natural properties—as the *fundamentalia*.

Lewis has certain hopes about what the *fundamentalia* will end up being like. Namely, he wants them to be local, intrinsic, and non-modal. They would be local if the fundamental entities occupied single spacetime points. The perfectly natural properties would be intrinsic if they were

⁴See Lewis (1973, 1986, 1994).

not somehow dependent on other entities and properties at other spacetime points.⁵ And finally, they would be non-modal if they were occurrent—not infected by things like dispositions, capabilities, or potentialities. There are significant questions about whether we should expect the fundamentalia delivered by physics to have all (or any) of these features. Certainly *current* physics, specifically quantum mechanics, can be seen as challenging at least some of Lewis’s desiderata.⁶ But these issues are not our primary concern here, so I will leave them aside.

The laws of nature do not appear anywhere in the Humean base. Rather, the laws are basically efficient *summaries* of what goes on in the base. More precisely: consider all of the particular facts comprising the Humean base. These may be organized into deductive systems, which allow us to derive certain facts from other facts. The laws, then, are the regularities of the deductive system whose predicates refer solely to the perfectly natural properties and that achieves the best balance of two desiderata: simplicity and strength.⁷ The strength of a system is a measure of how much information it provides about what goes on in the base. And the simplicity of a system is a measure of the simplicity of its axioms. The more informative a system is about what goes on in the base, the better. And the simpler it is, the better. Of course, these two desiderata often conflict. Great informativeness can be bought at the cost of simplicity (“there was a particle at spacetime point l_1 with mass m_1 and charge e_1 , and there was another particle at spacetime point l_2 ...”), and great simplicity can be bought at the cost of

⁵I do not mean to take any substantive stand here on precisely what it means for a property to be intrinsic, nor on which properties count as intrinsic and extrinsic. My present aim is only to summarize Lewis’s view.

⁶For discussion of the problems posed by quantum mechanics, see Maudlin (2007) and Loewer (1996), as well as some illuminating recent responses by Miller (2014) and Boghal and Perry (2017).

⁷Lewis says very little about what it takes for a statement to express a *regularity*. It had better be more than just a matter of the statement’s syntactic form, for any statement of particular fact can be forced into the form of a regularity. Again, though, we need not worry too much about the details here.

informativeness (“stuff happens”). Somewhere between these two extremes, at the optimal balance of both, are the laws of nature.

It is natural to present Lewis’s theory the way I have here because it mirrors the metaphysical structure of the laws, according to his view. The laws are not fundamental; they are “built” out of the fundamentalia. They are efficient summaries of the Humean base. So it is quite natural, in explicating his view, to start with a discussion of the Humean base and then proceed to discuss how the laws are derived from it.

However, in order to see the tension lurking in his view, we need to take a slightly different perspective. It will help if we ask where Lewis got this business about balancing simplicity and strength. He got it, he claims, from actual scientific practice:

...I take a suitable [deductive systematization] to be one that has the virtues we aspire to in our own theory-building, and that has them to the greatest extent possible given the way the world is. (1983, p. 367)

And again:

The standards of simplicity, of strength, and of balance between them are to be those that guide us in assessing the credibility of rival hypotheses as to what the laws are. (1986, p. 123)

The idea, then, is that scientists themselves try to balance simplicity and strength when they are investigating the laws of nature. So what Lewis is doing, essentially, is taking what we would traditionally regard as the *epistemic* principles for *figuring out* what the laws are, and elevating them to the status of *metaphysical* principles that are constitutive of what it is to *be* a law of nature (cf. Hall, ms, pp. 15-16). I will call the collection of these principles—which scientists use to investigate the laws of nature, and Lewis regards as *constitutive* of the laws—the “Nomic Formula.”

Of course, there is a significant question about whether Lewis was right about which principles figure in the Nomic Formula. Perhaps there are additional principles that scientists use when figuring out what the laws are, and perhaps they use things like simplicity in different ways than Lewis maintained (cf. Woodward (2014), Roberts (2008)). In fact I think both of these are the case, and in Chapter 2 I will argue for a revised conception of the Nomic Formula. But for now, I will not be overly concerned with details concerning the Nomic Formula's explication. What matters for now is the *relationship* between the Nomic Formula and the Humean base.

What is that relationship? According to the BSA, the laws of nature are the outputs of the Nomic Formula when it is applied to the Humean base. In other words, Lewis's theory is that if we take the totality of facts about the Humean base, and plug that into the Nomic Formula, whatever it outputs are the laws of nature. And of course, this is meant to be more than just extensionally adequate. According to the BSA, being a law of nature *just is* being an output of the Nomic Formula when it's applied to the Humean base.

2.2 Attractions of the BSA

Before we discuss the tension in Lewis's view, I want to consider what is attractive about it. Some of the most appealing features of the view can be traced to the elevation of epistemic principles to constitutive principles. Unfortunately, this is also one of the aspects of the view that Lewis himself talks the least about. Thus it is not entirely clear what his motivations were for making this move. Nonetheless, there are a number of things that make it appealing.

First, the BSA is attractive because it posits a relatively meager fundamental ontology. According to the BSA, the laws are not metaphysically fundamental; they are "built" out of the Humean base. When God is creating the universe, he does not have to first lay down the laws

and then the Humean base (nor vice versa); he just has to lay down the Humean base, and then the laws follow, as it were, for free. Thus, the BSA scores highly in terms of providing a parsimonious fundamental ontology. Insofar as ontological parsimony counts as a theoretical virtue, the BSA has it in spades.

Second, the BSA is attractive because it preserves the supervenience of the laws on the Humean base. Since being a law is a matter of being the output of a particular operation applied to the Humean base, it is not possible for two worlds to have the same Humean base but differ in their laws. Thus the view is often referred to as “Humean supervenience,” or “HS.”⁸ Now the natural question to raise at this point is: Why is it good to preserve supervenience? And the standard answer is that preserving supervenience is attractive for epistemological reasons. Specifically, imagine that the laws did *not* supervene on the Humean base. Then there would be a severe underdetermination worry: even if we could know all of the facts in the Humean base, we could not be sure that we knew the laws. In fact, it’s not only that we couldn’t be *sure*—it’s actually unclear how we could have *any* confidence that we had correctly identified the laws. For if supervenience were violated, then for any particular Humean base, there would presumably be an immense variety of different possibilities as to what the laws are. Violation of supervenience would thus leave us in an epistemological quagmire. But, the thought continues, we are in no such quagmire. We have very good reason to think, for example, that the Schrödinger equation, or some relativistic cousin of it, expresses a law of nature. Thus the laws must not violate supervenience. As Earman and Roberts (2005) put it: “the possibility of empirically justified belief, of a law of nature, that it is a law of nature, depends on HS” (p. 254).

⁸Strictly speaking, HS is the broader thesis that *all* truths, not just the truths about the laws, supervene on the Humean base.

Now, supervenience is perhaps a necessary condition for the epistemic accessibility of the laws, but it is not a sufficient condition. Imagine, for example, an account according to which the laws are an uncomputable function of the facts comprising the Humean base. Such an account preserves the laws' supervenience, but it does not preserve their epistemic accessibility. What is important to note, then, is that the BSA does more than just preserve supervenience: it makes the laws epistemically accessible. For according to the BSA, the laws count as laws in virtue of satisfying the *very criteria* that scientists actually use to discover them.⁹

This naturally leads to the third, and most significant, attraction of the BSA. By elevating scientists' epistemic standards to constitutive standards, we are able to make the *most possible sense* out of scientists' epistemic practices. In other words, it is as if Lewis asked himself: "How should one design a theory of laws to make sense of the ways that scientists investigate them?" And his response was this: "Figure out what sorts of things scientists' epistemic practices would be good at discovering, and make those things the laws." Now, regardless of whether Lewis correctly articulates the particular epistemic standards that scientists use to investigate the laws, this strikes me as a sound methodological maneuver. By employing it, we automatically avoid a variety of skeptical possibilities, whereby scientists might turn out to be using investigative principles that would systematically mislead them about what the laws of nature are. We also provide an immediate justification for why scientists should be using those very investigative principles: scientists search for principles that satisfy the standards in the Nomic Formula because satisfying those standards is just *what it is to be a law*.¹⁰

⁹At least, the BSA is traditionally regarded as attractive for this reason. In §3, I will articulate some doubts about the claim that the BSA renders the laws epistemically accessible.

¹⁰These considerations also undermine the otherwise natural concern that the epistemic-to-constitutive elevation is just a conflation of metaphysics with epistemology. Essentially, the epistemic-to-constitutive elevation is justified by a prior commitment to the rationality of our scientific investigations into the laws. That, to me, seems a very compelling reason for letting epistemology be a guide to metaphysics.

To make this point clearer, let's suppose we *reject* the constitutivization of scientists' epistemic principles, and see how this generates a philosophical problem. Now, regardless of whether the epistemic-to-constitutive elevation is correct, it is undeniably part of the task of science to discover the laws. However, if we reject this elevation, we are faced with a troubling fact about how scientists actually go about investigating the laws. Lewis was right, after all, that they tend to value simpler theories over more complex ones, other things being equal.¹¹ But if satisfaction of our epistemic principles is *not* part of what it is to be a law, then why should those principles, specifically the preference for simplicity, be a guide to the facts about the laws? That is, why should we have higher confidence in simpler theories about the laws than in more complex theories?

Essentially, by rejecting the epistemic-to-constitutive elevation, we are introducing a wedge between our epistemic standards for discovering the laws and the laws themselves. We are thus faced with the task of explaining why those standards are a reasonable guide to the laws. Specifically, we are faced with the task of giving an epistemic justification of simplicity vis-à-vis the laws of nature. And this has turned out to be a very difficult task indeed. We have no satisfactory general account of why simplicity should be a guide to truth, nor why it should be a guide to truth with respect to the laws of nature. Lamenting our lack of success in finding such an account, J. J. C. Smart said that our expectation of simplicity in nature perhaps “derives from earlier theological notions: we expect God to have created a beautiful universe” (1984, p. 121). Of course, in certain contexts, with certain background assumptions or desiderata, it is possible

¹¹Again, I am not assuming that Lewis was right about exactly how scientists value simplicity in their theories. But I *am* assuming (rather innocently, in my opinion) that they do value it. I address the role of simplicity more thoroughly in Chapter 2.

to give limited epistemic justifications of a preference for simplicity.¹² Woodward (2014) presents several such justifications, but the critical point is that these are both context-sensitive and goal-dependent. Thus they do not go far toward establishing the general reliability of simplicity with respect to theories about the laws of nature. Woodward himself makes this point in a *critique* of the BSA:

Lewis and many of the other philosophers who have defended the BSA have [...] little that is illuminating to say about why we should prefer simpler hypotheses, in the sense of providing some argument that such a preference leads to the successful identification of laws [...] Instead, defenses of the BSA tend to treat the preference for simplicity as well as the notion of simplicity itself as a kind of primitive, and build these into the characterization of laws (laws just are generalizations that figure in systematizations best combining simplicity and strength) so that questions about the relationship between simplicity and lawfulness, and why we should value simplicity are not given non-trivial answers. It would be much more satisfying if defenders of the BSA [...] were to provide some reason to suppose that trading [simplicity] off against strength tracks the features that we antecedently think are possessed by laws. (2014, pp. 107-108)

Here Woodward gives an eloquent characterization of one of the most attractive features of the BSA, and then immediately turns around and calls it unsatisfying. The BSA is attractive *precisely because* it lets us avoid the task of showing why simpler hypotheses about the laws are likely to lead to the “successful identification” of the laws. There’s no need to try to connect the dots when there is *only one dot*. But Woodward says that it would be more satisfying if the BSA *did* try to connect the dots, i.e. if it gave us some reason to think that simplicity will accurately track the laws of nature. Now, it certainly isn’t my place to tell Woodward what he should and should not find philosophically satisfying. I would only remark that it is unfair of him to level this objection solely against the BSA. For scientists *do* use simplicity in their investigations about the laws of nature, so *any* theory of laws has to be able to make sense of this. The BSA just has a particularly easy way of doing so. (This is what Hall [ms., p. 38] calls a “nifty judo move”

¹²See, for example, Sober’s (1994) excellent discussion of simplicity considerations in the construction of phylogenetic trees.

that reductionists like Lewis can perform on their non-reductionist opponents.) Therefore Woodward's objection apparently boils down to the claim the BSA gets off easy where other theories of law have encountered an intractable problem. This seems an odd reason to *dislike* a theory.

Once we stop thinking of simplicity as solely a *guide* to the laws, and instead think of it as partly *constitutive* of being a law, then we are free to reinterpret its justification. Specifically, if simplicity helps to constitute the laws, then it is *trivial* why we should think of simplicity as a guide to the laws, and the question now becomes why we should be *interested* in discovering a set of truths that is partly distinguished based on its simplicity. In other words, we are free to search for *pragmatic* justifications of the preference for simplicity. And these are easy enough to come by: simpler hypotheses are easier to work with and memorize, they make calculations more efficient, and so on.¹³

It is worth reiterating that I think Lewis left out some other aspects of the Nomic Formula—that is, standards that scientists try to balance aside from just simplicity and strength. For example, in Chapter 2 I will claim that scientists tend to look for principles that have very wide applicability. Thus, by constitutivizing scientists' epistemic principles, we make wide applicability part of what it is to be a law, and we thereby make sense of why scientists look for such principles when they are trying to discover the laws. Conversely, by *not* constitutivizing their epistemic principles, we make it an *utter mystery* why scientists look for principles with wide applicability when they are trying to discover the laws. Thus, by developing a more accurate and detailed account of the Nomic Formula, we can strengthen the case against theorists

¹³There is much more to say about the justification of simplicity as part of the Nomic Formula. One pertinent question is why it is pragmatic to balance simplicity against the other desiderata in the formula. I address this in the next chapter.

like Woodward who want some non-trivial reason to think that the elements of the Nomic Formula will accurately track the laws of nature. Once we see all the elements of the Nomic Formula in detail, the burden of *not* constitutivizing our epistemic principles becomes more severe, for then we would need to explain, not only why simplicity is a good guide to the laws, but also why a variety of other standards are as well.

In sum, there are a number of reasons that Lewis's elevation of investigative principles to constitutive principles is quite attractive. Loewer (2007) gives an apt summary of the nice results of this move:

[The BSA is attractive] because of the way it incorporates the criteria physicists use for counting generalizations and equations as expressing laws and also because it, unlike many of its rivals, doesn't posit metaphysically primitive laws, primitive causal powers, propensities, governing relations, or other metaphysically heavy-duty and suspect entities. (p. 313)

But now, given that the BSA is so attractive, you might wonder why we don't perform this epistemic-to-constitutive elevation all over the place. That is, why don't we say, for example, that *planets* are the outputs of our planet-investigation procedures? Or why don't we say that *persons* are the outputs of our person-discovery principles? For clearly we *don't* say anything so ludicrous; the move is not even slightly attractive in these other cases. So what makes the epistemic-to-constitutive elevation legitimate in some cases but not in others? And why is it legitimate in the case of laws?

I doubt that there is a definitive, a priori way to know when the epistemic-to-constitutive elevation will be an appropriate maneuver in the construction of a philosophical theory. The ultimate test will be how strong of a theory we can create by employing it. I hope to convince the reader that a very strong theory can be created by using the epistemic-to-constitutive elevation in the case of laws. Presumably, a parallel theory of planets or persons could *not* be made very

plausible. It is an interesting question why that is the case, but not one that I feel particularly obliged to answer here, since this is not a philosophical treatise on either planets or persons. I would suggest, however, that it may have something to do with the fact that, in the cases of planets and persons, there are fairly obvious *objects* we can point to in their identification. The availability of such objects obviates the need for some *alternative* way of identifying the entity in question. By contrast, in the case of laws, there is nothing we can point to, in any straightforward way, and say, “That thing is a law of nature.” Therefore, some other identifying criteria are necessary, and that is exactly what our epistemic standards for investigating the laws provide.¹⁴

Before wrapping up this section, I want to address a natural objection to raise at this point.¹⁵ What Lewis and I are recommending is that we regard our principles for investigating the laws as principles that are constitutive of being a law. But one might think that this is taking an overly narrow view of things. The laws have a *variety* of distinctive features. Here are several:

1. Laws support counterfactuals
2. Laws have a characteristics necessity, but they are also contingent in some respect
3. Laws are important features of our world worth knowing
4. Laws are things that scientists try to discover using the relevant epistemic principles
(whatever those may be)
5. Laws are useful for predictive purposes
6. Laws underwrite reliable measurement procedures
7. Laws explain natural phenomena

¹⁴One other thought: perhaps the reason that it sounds implausible to say that planets are the outputs of the ideal implementation of our planet discovery principles is that it has the flavor of something quite uninformative. (Recall Wittgenstein’s, “By ‘brick!’ I mean ‘brick!’” (1965).) Perhaps if those discovery principles were spelled out in more detail, the account would begin to sound more plausible.

¹⁵Thanks to Marc Lange for pressing this objection.

8. Laws are confirmed by their instances
9. Laws govern the evolution of events in the universe

Some of these features overlap. For example, the laws characteristic necessity might be thought to connect up with their counterfactual support: since they are necessary, they are held fixed in the evaluation of counterfactuals (or, perhaps, since they are held fixed in the evaluation of counterfactuals, they are necessary). Now, my account focuses explicitly on feature (4).¹⁶ I have already argued that this constitutes a plausible starting point: by taking scientists' epistemic standards as central to our account of laws, we are able to make the most possible sense out of those standards. Other accounts of laws in the literature have focused on other features. For example, Roberts (2008) focuses largely on feature (6); Lange (2000, 2009) focuses most directly on (1); and Maudlin (2007) focuses on (9). I cannot offer any argument that all starting points other than my own are guaranteed to result in inferior theories. My hope is only to demonstrate that, by taking scientists' epistemic standards as the foundation for a theory of laws, we are able to provide a very illuminating, very powerful account of a variety of features possessed by the laws, and of a variety of ways that the laws function in scientific practice.

Now, you might think that all of these theories constrain themselves in unnecessary ways. That is, they all take some aspect of the laws as a starting point, and try to use it to make sense of as many other aspects as possible. Call this the "starting point" methodology. Another approach would be to not take any one of these features as central to lawhood, but instead use them all to *jointly* fix what lawhood must be. Roughly, we enumerate the various roles played by the laws, and then we ask: "What sort of thing could play as many of these roles as well as possible?" Such a "conceptual role" methodology has been suggested by people like Beebe (2000) and

¹⁶In Chapter 2 I will argue that features (3) and (5) are in fact natural correlates of (4), so it would be appropriate to say that my starting point is really the cluster of features (3), (4), and (5).

Loewer (1996) in developing accounts of laws. One attraction of this methodology is that it appears to give us the best chance of making the most sense of *all* of these features, as opposed to just the one that we start with. Thus I am confronted with the following question: Why take any *particular* feature of the laws to be our starting point? Why not instead adopt a conceptual role methodology?

I do not think there is any *a priori* reason to think that such a methodology would fail. The ultimate test of a philosophical methodology is the results that it generates, and indeed, it is not inconceivable that one could use a conceptual role methodology to develop a theory of laws that is very similar in spirit to my own. To be sure, I very much hope that the laws of my theory *do* fare, on balance, better than all others at playing these various roles. But to say this is only to commit myself to the claim that conceptual role considerations serve as a suitable test for a *completed* theory. It does not commit me to the logically independent claim that conceptual role considerations are the best guide to the *development* of such a theory. So while I am in favor of conceptual role considerations at the evaluation stage, I have some reservations about explicitly trying to develop a theory of laws along these lines.

Oftentimes, I think that the use of a conceptual role methodology only serves to disguise the true motivations behind a theory, in the following way. The theorist asks themselves, “What sort of thing could play all or most of these roles reasonably well?” But what they are implicitly doing is thinking of some of these roles as more important, or more significant, than others. These tacitly-more-significant roles may vary from theorist to theorist, depending either on explicitly-reasoned arguments, or on pre-theoretic dispositions and intuitions. Reductionists, for example, tend to privilege the epistemic features of the laws, e.g. the fact that we are in a position to be able to discover them. Nonreductionists, on the other hand, tend to privilege things

like the laws' characteristic necessity, or their governing relation over the sub-nomic facts. So both sides will design theories that respect their privileged features, and the challenge is then to find a way to address the features that have been implicitly discounted. Usually this is done by providing additional philosophical theories. So, for example, by taking the laws to be a certain sort of pattern in the phenomena, the reductionist tends to have trouble accounting for the fact that the laws support counterfactuals. (If laws are mere patterns, why should they be held fixed in the evaluation of counterfactuals?¹⁷) What the reductionist needs is then a theory of counterfactuals that is (a) tailored to their theory of laws, and (b) able to account for the laws' counterfactual support in its own terms. For an example of this dialectic in practice, see Loewer's (1996). There he articulates a number of roles the laws are supposed to play (some of which I borrowed from him in my list above), and evaluates how well Lewisian laws ("L-laws") can play these roles. He argues that L-laws can play the roles required of them reasonably well if we understand those roles in Lewis-friendly ways (i.e. in terms of "L-counterfactuals," "L-necessity," and "L-explanation," to use his terminology).

Now, what is the difference between this methodology, and one on which we explicitly start with some privileged feature of lawhood, and attempt to salvage as many other features as possible, constructing new theories of counterfactuals, necessity, explanation, etc. as required? To my eyes, the difference is merely superficial; it lies in how the approaches are *presented*, not in how they actually proceed. In other words, the primary difference (in my opinion) is that theorists who explicitly proceed using the conceptual role methodology are merely being less up-front about their true starting point. Of course, this may not be a conscious decision on their part; such theorists may not realize that their approach weights certain conceptual roles more heavily

¹⁷I develop an answer to this question, consonant with my theory of laws, in Chapter 3.

than others. But, come to think of it, *shouldn't* some roles be considered more important than others? Surely, for example, the fact that the laws bear a particularly tight relation to counterfactuals is a more important feature to account for than the fact that the laws are referred to by English speakers using the four-letter word “laws.” The former feature could plausibly be placed at the center of an tenable account of laws, but an account that placed the latter feature at its center would be doomed to failure.

Once we admit that certain conceptual roles are more important than others, we have departed somewhat from the guiding idea behind the conceptual role methodology (at least as I have characterized it here). For the idea of a “starting point” now finds a bit of purchase: the starting point for our theory of laws is the role (or cluster of roles) that we take to be the most important. How would we then proceed to develop a theory? Presumably, we design our theory of laws with that role at the center, and our task is to account for as many other roles and features of the laws as possible, construed as natural extensions of that starting point. And this is just the “starting point” methodology I described above. It is a methodology I think *all* philosophical theories are bound to in one way or another.

This completes my review of the attractive features of Lewis’s BSA, and my justification for developing a theory in this manner. In short, there is a great deal to like about this approach. Unfortunately, there is also a serious tension at the heart of it.

3. A Tension in the BSA

The tension I am interested in is centered on Lewis’s characterization of the Humean base. For example, Barry Loewer (2007) argues that Lewis’s notion of naturalness, which determines which properties are part of the Humean base, introduces an epistemic divide

between (1) all of our possible evidence relevant to the laws of nature, and (2) the laws themselves. The thought, which he attributes to van Fraassen (1989), is the following. First, note that we do not seem to have any definitive way of telling whether our predicates and kind terms refer to perfectly natural properties. With this in mind, imagine that scientists arrive at a fundamental physical theory that they *think* counts as the Best System. This theory reaches the best balance of all the investigative principles that scientists are able to evaluate (i.e. simplicity, strength, and whatever else figures into the Nomic Formula). And it contains certain putative law statements. Nevertheless, it is possible that these are not the actual laws, because it is possible that the theory does not refer exclusively to the perfectly natural properties. The *real* Best System has to do so, and thus the real laws might diverge from what we have every reason to believe are the laws. As Loewer puts it,

[t]he epistemological problem is that on Lewis's account even knowing all the non-nomological contingent truths [...] isn't sufficient for knowing which truths are the laws. One would also have to know which predicates refer to Lewisian natural properties. (2007, p. 322)

There are several potential replies open to Lewis, but it would take us too far afield to review them all here; for such a review, see Loewer (2007) and Cohen and Callendar (2009), both of whom argue forcefully that none of these replies work. What is important for now is the general tenor of this worry: Lewisian naturalness, and by extension the perfectly natural properties that figure in the Humean base, may be epistemically inaccessible.

Cohen and Callendar (2009) raise a related worry. Here is what they say:

The Best System approach, by its very nature, introduces a metaphysical asymmetry where we have an epistemological symmetry. That is, scientists actually devise laws based on their choice of kinds and choose their kinds based on the laws. Gell-Man[n] was not simply handed fractional charge and left to make the best system he could with it;

rather he postulated fractional charge in part because he saw that he could make a very simple and strong system with it if he did.¹⁸ (2009, p. 13)

The thought here is that the BSA drastically misrepresents how actual scientific practice works. Scientific practice does *not* proceed like this: we are *given* the fundamental magnitudes and the behaviors of the fundamental entities, and out of careful study of these phenomena, we derive some principles that are the laws of nature. Rather, it proceeds like this: both the laws and the fundamentalia are derived *together* to account for *other* data. We will have to consider later (§4) just what these “other data” are, but the important point at present is that the BSA misleadingly suggests that the fundamentalia are part of the inputs to our theorizing about the laws. Instead, they are more appropriately regarded as part of the outputs.¹⁹

With these objections on the table, the tension in the view is now readily apparent. On the one hand, Lewis thinks our theory of laws ought to respect the epistemic principles that scientists use to *investigate* the laws, i.e. the Nomic Formula. On the other hand, Lewis thinks our theory must involve application of the Nomic Formula to the entities that are ontologically fundamental: particles, fields, strings, or whatever our best physics tells us. But *scientists* never apply the Nomic Formula to these kinds of entities. In actual scientific practice, the fundamental entities and properties are just as much part of the outputs of the Nomic Formula as are the laws

¹⁸I cannot help but remark on how aptly Cohen and Callender’s claim summarizes Gell-Mann’s original motivations. In the 1964 paper in which he introduced the idea of the quark, Gell-Mann proposed two alternative explanations of the “eightfold way,” a highly successful classification system for hadrons. Both of these explanations involved the idea that hadrons were composed of more fundamental particles, but the first of them introduced a somewhat inelegant asymmetry in the posited properties of these more fundamental particles. In light of this, Gell-Mann suggested, “A simpler and more elegant scheme can be constructed if we allow for non-integral values of the charges” (1964, p. 214).

¹⁹Indeed, in some cases laws are used to establish the existence of new fundamentalia. For example, certain conservation laws essentially have a foundational status in particle physics. Many detectors used in particle colliders are unable to register uncharged particles. However, given the energy and momentum conservation laws, researchers infer the existence of uncharged particles as the result of imbalances in pre- and post-collision energy and momentum. Any decreases detected in energy-momentum between the interacting and resulting particles must be accounted for by positing that the collision also resulted in uncharged (and therefore undetected) particles.

themselves. In a word, then, the tension is this: the BSA directs us to apply the Nomic Formula, which is derived from scientific practice, to data that *scientists* do not apply it to, to which it was *never intended* to be applied.

So what I am suggesting is that there is a motivational question about the basic architecture of Lewis's view. Lewis says we are to apply the Nomic Formula to the Humean base, but given his explications of those two elements, it is puzzling why we should expect this recipe to generate the laws. Additionally, quite apart from this motivational question, the tension between the Nomic Formula and the Humean base also generates some problematic downstream consequences. In particular, consider the following two possibilities:

1. Physicists reach the conclusion that at least some of the fundamental entities are in principle undetectable.
2. Physicists reach the conclusion that there is no fundamental level.

Both of these cases, I take it, are at least possible. It is instructive to ask about the consequences for the BSA in each of them.

Consider case (1). There are presumably a variety of ways physicists might reach the conclusion that the fundamental entities are in principle undetectable. This may involve the presence of certain asymmetries in the properties and laws governing *detectable* entities. Positing undetectable entities may smooth over these asymmetries in a way that makes the overall theory simpler, more elegant, etc. Or it may involve certain symmetries in the detectable entities that would be nicely explained by another ontological level. The most obvious example from contemporary physics is that of quarks, which have never been detected in isolation. This is due to a phenomenon called “confinement,” according to which color-charged particles (like quarks) cannot exist by themselves. The so-called “virtual particles” of quantum field theory constitute

another contemporary example. These particles do not obey Einstein's energy-momentum relationship.²⁰ They exist for extremely small time intervals, making them experimentally undetectable.

In such a case, it is not only that scientists *in fact do not* use data about the behavior of these fundamental entities in their theorizing about the laws. Rather, scientists are *excluded in principle* from using such data. The BSA would thus direct them to apply the Nomic Formula to data to which they cannot apply it. In fact, to make the problem starker, it will help to imagine that—per impossibile—scientists have *all* the data about the detectable phenomena, but they lack data about the undetectable, fundamental phenomena. In such a case, the BSA will say that scientists may be wrong about the laws, though they will never be able to know it. The pressing question for Lewis at this point would be: Why does your theory recommend the possibility of this in-principle-undiscoverable error?

Now, it might be thought that this possibility is relatively unobjectionable. After all, we are accustomed to the fact that scientists can be incorrect in their claims about the laws of nature. They can, and do, make errors. So maybe Lewis can maintain that his theory is revealing what the laws *really* are, not just what scientists might incorrectly conclude about them in the case where the fundamentalia are undetectable. And of course, Lewis *can* maintain this. But then what we ought to ask is this: On a reductive theory of laws like Lewis's, what is the *motivation* for doing so? In other words, why would it *make sense* for Lewis to allow that the laws may be underdetermined by all of our possible evidence?

I certainly understand countenancing the possibility—or rather, the *necessity*—that the laws are underdetermined by all of our *current* evidence. For we can always observe something

²⁰ $E^2 = p^2c^2 + m^2c^4$, where E is the particle's energy, p the momentum, m the mass, and c the speed of light. Particles that violate this relationship are said to be "off mass shell."

tomorrow that conflicts with our best theories about the laws today. I even understand countenancing the possibility that the laws are underdetermined by all the evidence we will ever actually gather. For if we didn't countenance such a possibility in our theory of laws, then it would be very difficult to explain why scientists go to such great lengths to gather additional data in order to discover the laws. But why countenance the possibility of underdetermination by *all possible* evidence? For a view whose primary attractions are that it (a) renders the laws epistemically accessible, and (b) makes sense of scientists' epistemic standards, it seems gratuitous to allow for the possibility that we could be wrong about the laws even if we had gathered all possible evidence and perfectly applied our epistemic standards to that evidence.

Of course, someone who thought that the laws were entities with an existence independent of the particular matters of fact could easily ground the underdetermination of the laws by all possible evidence. Even if we had all possible evidence, and even if we reasoned perfectly from that evidence to a theory of laws, that theory could still be false, because the ultimate test of a theory of laws on such a view is simply this: Do the laws, these independently-existing entities, match our theory's claims about them? But this is not the way a reductionist like Lewis thinks about the laws; for him, the laws manifestly *are not* independently-existing entities. So on a reductionist view of laws, it is not clear what motivation there is for maintaining that we *could* be wrong about the laws in such a case. All that Lewis seems to be able to say is that *his* theory allows that we can be wrong about the laws when we have gathered all possible evidence, and reasoned perfectly about that evidence. So much the worse, we might think, for his theory.

Now, I am not suggesting that this is a *damning* argument against Lewis. But at the very least, it should be granted that in the situation described by case (1), the BSA mandates a

divergence from scientific practice in the following sense: it tells us to apply the Nomic Formula to data that scientists do not—and cannot—apply it to. This divergence opens up the possibility of an in-principle-undiscoverable error on the part of scientists. Thus Lewis at least faces the task of motivating that possibility. My suggestion is that it is very difficult to motivate that divergence from within a reductionist theory of laws like Lewis's, whose attractions are largely epistemological.

Now consider case (2). Here we are supposing that fundamental physicists reach the conclusion that there is no fundamental level. To use an already-over-used metaphor, it is “turtles all the way down.” Such a possibility is perhaps a bit unintuitive, but it is not unheard of. For example, Ladyman and Ross (2007) entertain this as a “tentative hypothesis.” They maintain that “[w]e have inductive grounds for denying that there is a fundamental level since every time one has been posited, it has turned out not to be fundamental after all” (p. 178). Schaffer (2003) also argues that the assumption of a fundamental level is one about which we should be agnostic. In fact, he suggests that the assumption plays a *pragmatic* role in the funding and practice of physical research, and it is for this reason, rather than its overwhelming plausibility, that the assumption enjoys such widespread acceptance.

If there were no fundamental level, what would be the consequences for Lewis's view of laws? Well, again, according to Lewis the laws are the results of applying the Nomic Formula to the Humean base, i.e. to the totality of facts about the fundamental level. So if there *is no* fundamental level, it follows straightforwardly that there are no laws. The BSA thus makes the existence of laws of nature dependent on the existence of a fundamental level.

This is an odd conclusion. Contemplating the possibility that there is no fundamental level to reality may indeed produce a sort of vertigo, but it is distinct from the “inductive

vertigo” one gets from supposing that there may be no laws of nature.²¹ Neither Schaffer nor Ladyman and Ross make any suggestion that the absence of a fundamental level rules out the existence of laws of nature. Furthermore, no physicist that I know of who has contemplated the possibility that there is no fundamental level of reality has suggested that this would preclude the possibility of any laws of nature. David Bohm (1957), for example, suggested the possibility of “an infinity of levels” in the physical world (p. 138). But he did not suggest that such a hypothesis undermines any laws of nature, and in fact he himself proposed a novel law of nature in the form of the “guiding equation” of the Bohmian interpretation of quantum mechanics. Bohm would thus have been one of the last people to suggest that there are no laws of nature. Another example is Geoffrey Chew’s (1962) “bootstrap model” of strong interactions, according to which there are no truly fundamental particles. The bootstrap model became prominent in particle physics in the 1960s, and was motivated by the rapid proliferation of officially recognized hadron particles as the result of developments in linear accelerator technologies. Advocates of the bootstrap model attempted to do away with the idea that some particles are more fundamental than others; rather, there are just different sorts of particles created in different conditions. However, far from denying the existence of laws of nature, they were attempting to understand the nature of the strong interaction, i.e. the mechanism responsible for the strong nuclear force.

Furthermore, many putative laws of nature would seem to be laws *regardless* of whether there is a fundamental level or not. Take, for example, Newton’s laws of motion. These are formulated quite generally, without referencing any fundamental kinds of entities or properties. They refer to things like forces, velocities, masses, and accelerations, but these can be properties

²¹I borrow this term from Blackburn (1993, p. 98).

of *any* physical system. So the lawhood of Newton's laws does not depend on the existence of a fundamental level.²² Similarly, the Schrödinger equation would appear to still be a law regardless of whether there is a fundamental level. Consider the form of the equation:

$$i\hbar \frac{\partial \Psi}{\partial t} = H\Psi$$

None of the terms here depends on the existence of fundamentalia. Of course, i is just a mathematical term and does not denote any physical quantity. The term Ψ denotes the state vector of a quantum system, and in principle *any* physical system can be associated with a state vector—an electron, a Buckminsterfullerene molecule, the pencil resting on my desk, and even the universe as a whole. Similarly, H represents the Hamiltonian of the system, i.e. the total energy. It is calculated differently depending on the system in question, but again, none of its components must be fundamental quantities. Finally, \hbar is the reduced Planck constant, which is Planck's constant (h) divided by 2π . It has the best claim to being a fundamental quantity here, for h is the “quantum of action”—it relates the energy contained in a photon with the frequency of that photon. It is a fundamental *constant*. But just like the gravitational constant does not necessitate a fundamental level, neither does Planck's constant. There can be fundamental constants in the universe without there being a fundamental level.

Richard Feynman has likewise emphasized that most laws are, in a sense, subject neutral. They do not specify the sorts of entities that they apply to. Throughout his series of lectures on *The Character of Physical Law*, he discusses various laws of nature. At the beginning of the last lecture, then, he says the following:

You may think I have told you everything already, because in the lectures I have told you all the great principles that are known. But the principles must be principles about *something*; the principle of the conservation of energy relates the energy of *something*,

²²Granted, Newton's laws are *false*, and thus strictly speaking they are not *laws*, but the reason they are not laws has nothing to do with the fact that they do not refer to fundamentalia.

and the quantum mechanical laws are quantum mechanical laws about *something*—and all these principles added together still do not tell us what the content is of the nature that we are talking about. I will tell you a little, then, about the stuff on which all of these principles are supposed to have been working. (1965, p. 149, italics in original)

In short, the issue of which entities the laws operate on is distinct from, and cannot be read off of, the laws themselves. One can discuss the laws without discussing those entities, fundamental or otherwise. Thus the laws do not require or presuppose the existence of fundamentalia.

Now, there is a natural objection to raise at this point. What we are most concerned with, after all, are *fundamental* laws of nature. So according to the objection, whereas derivative or non-fundamental laws may not require the existence of fundamentalia, fundamental laws of nature surely *do* require this. After all, it is plausible to suppose, as Lewis does, that the fundamental laws are the laws governing the fundamentalia. What makes them fundamental is that their relata are themselves fundamental. On this understanding of what it is to be a fundamental law of nature, if there is no fundamental level, there can be no fundamental laws. So Lewis is vindicated.

The foregoing discussion, however, suggests otherwise. Laws of nature—fundamental or otherwise—are not held hostage to the existence of a fundamental level. There could be fundamental laws without a fundamental level. What is meant by calling some laws “fundamental” is that they are *explanatorily prior* to all other laws. And the explanatory priority of a law need not track the fundamentality of its relata. For example, Lange (2007) has emphasized that it may be no coincidence that all of the force laws are spatially and temporally symmetric. (That is, the force laws are invariant under spatial and temporal transformations.) Rather, these “symmetry principles” might constrain the way the force laws could be; if the particular force laws had been different, they would still have exhibited temporal and spatial symmetry, because the symmetry principles would still have held. Thus, symmetry principles

would be explanatorily prior to the particular force laws. But notice that there is nothing in the statement of the symmetry principles that refers to or depends on the existence of fundamental entities. They state merely that all force laws are invariant under spatial and temporal displacement. (In fact, it is not entirely obvious exactly what the relata are in this case. Perhaps it could be said that the relata are the individual force laws.) Thus symmetry principles may be more fundamental than the particular force laws, but this is not in virtue of the relative fundamentality of their relata. It is in virtue of their explanatory priority.²³

This conclusion is in keeping with the aforementioned discussions of physicists and philosophers regarding the possibility of the lack of a fundamental level: none of them suggested that it precludes fundamental laws of nature. So I stand by my previous claim that there may still be fundamental laws of nature even if there is no fundamental level. Lewis is un-vindicated.²⁴

In sum, as the term “law of nature” functions in scientific practice, it appears to have no tight connection with the existence of a fundamental level. It would therefore be quite odd to impose such a connection within a philosophical theory of laws of nature. But this is exactly what the BSA does. The laws are the outputs of the Nomic Formula when it is applied to the totality of fundamentalia, so if there are no fundamentalia, there are no laws. This underscores the fact that in actual scientific practice, the laws are not discovered by applying the Nomic

²³I will have more to say about these symmetry principles in my account of the Nomic Formula in Chapter 2.

²⁴One other point deserves mention here. John Roberts has stressed to me that, in the official statements of Lewis’s view, he never explicitly says that the Humean base consists of the totality of facts about fundamentalia. Rather, it consists in the totality of facts about the spatiotemporal distribution of local qualities. So it may be suggested that Lewis’s view is not problematized by the possibility that there is no fundamental level. As long as the distribution of local qualities exists (at some level or other), Lewis would still have the supervenience base that his view requires.

Though Lewis doesn’t explicitly say that the Humean base consists of the fundamental facts, I take it that this is clearly what he intended. In his (1994), he says that defending Humean Supervenience is valuable because, even if physics tells us there are no local qualities, “that defence can doubtless be adapted to whatever supervenience thesis may emerge from better physics” (p. 474). So Lewis seems to be thinking that if the fundamental physical properties do not meet his stipulation of spatiotemporal locality, the laws still supervene on the distribution of whatever non-local facts constitute the fundamental level. Thus I think that the possibility of the lack of a fundamental level strikes at the underlying idea of his view, if not at the official statement.

Formula to the fundamentalia. As we have seen, the fact that the BSA says otherwise generates some very odd consequences for it.

We might then wonder why Lewis structured his theory the way that he did.

Unfortunately, he is pretty quiet about his motivations on this front. I am not sure why he thought that applying the Nomic Formula to the fundamentalia would give us the laws of nature. One gets the sense that he might have been trying to accord *too much* respect to actual physical practice. What I mean is this. There are at least two ways that a philosophical theory can pay heed to actual physical practice. On the one hand, a philosophical theory can make use of *results* from actual physics, and on the other hand a philosophical theory can make use of physical *methods*.²⁵ I am of course being intentionally vague about what it means for a philosophical theory to “make use of” the results or methods of physics. There are a variety of ways to do both. Importantly, however, there are certain permutations of these ways that do not make sense to combine into a single philosophical theory. My overall suggestion is that the BSA attempts to use both the results and the methods in a way that does not make sense. In actual physical practice, the methods are used to arrive at the results. But Lewis wants us to apply the methods, i.e. the Nomic Formula, *to* the results, i.e. the fundamentalia. This is odd, for the results were only reached in the first place by the application of the method. It’s as though we’re being told to apply the method twice. *That* is according too much respect to physics—more than physics accords itself!

Hopefully by now I have convinced you that the tension inherent in the BSA is problematic. Of course, I do not mean to suggest that it is *disastrous*. Lewis could very well accept these odd consequences and hold onto the BSA in its present form. As I see it, the main

²⁵Cf. Russell (1914).

motivation for doing so would be to preserve the appealing features I mentioned back in §2.2. It might be suggested that yes, ultimately there is a strange tension at the heart of the theory, but on the other hand it has all of these other wonderful benefits: it makes sense of the scientist's appeal to simplicity (as well as their other epistemic standards), it renders the laws epistemically accessible, and so on. Fortunately, however, I think we can eliminate the tension and maintain those benefits.

4. Reconceptualizing the Humean Base

Given that it does not make sense to apply the Nomic Formula to the Humean base, we have two options: we can modify our conception of the Nomic Formula, or we can modify our conception of the Humean base. (I suppose we also have a third option: throw it all out. But that would be premature. The general architecture of Lewis's theory is not itself defective; the problem is that the particular pieces do not fit together.) Now, obviously I am going to argue that we should reconceptualize the Humean base. But why should we not instead change the Nomic Formula?

Well, in a sense I think we *do* need to change the Nomic Formula also. I have already mentioned that I think Lewis left out some of the desiderata that scientists attend to when reasoning about the laws, and those should be added in. But putting that aside, the general thought behind the Nomic Formula is sound: we examine scientific practice, and extract the principles that scientists use in their reasoning about what the laws are. Being a law of nature, then, is a matter of being the output of the (ideal, mistake-free) operation of these principles. I have already argued at length for this conception of the Nomic Formula back in §2, so I will not recapitulate here.

If we stick to this overall conception of the Nomic Formula, then in light of its tension with the Humean base, we ought to consider how to revise our conception of the latter. The problem with Lewis's conception was that it was too divorced from the sort of data to which scientists actually apply the Nomic Formula. Lewis attempted to take something like a God's-eye view when characterizing the Humean base. But scientists do not—*cannot*—take that view, and consequently the Nomic Formula is not designed to be applied to a world described from such a view. The natural reaction, then, is to align the Humean base more closely with the data to which scientists actually apply the Nomic Formula. In other words, my suggestion is that we should regard the Humean base as the totality of data that scientific theorizing aims to systematize and explain. In a word, the Humean base is the totality of “input data” to scientific theorizing.

Of course, just saying that the Humean base is the “input data” does not provide a very good sense of what that sort of data is. Unfortunately, it turns out to be very difficult to explicate, in a precise and general way, just what is meant by the “input data,” and philosophers of various persuasions are likely to disagree significantly about exactly how it should be characterized. However, two points here work in my favor. First, it will not matter too much *exactly* how we characterize the “input data,” for the overall picture offered by my view will not be drastically affected by such details. Thus in my final analysis, I will leave certain aspects of the characterization of the input data unspecified; these can be filled in by philosophers of various persuasions depending on their own sympathies. Second, it is important to note that this is not just my problem—it is *everyone's* problem. The question of how to characterize the sort of data that scientific theorizing aims to systematize and explain is one that has to be answered regardless of what further theoretical uses we might put that answer to. So my theory of laws would not have to be scrapped if it were found to rest on an inadequate analysis of the input data.

Rather, we would only need to revise my analysis of the Humean base to bring it in line with the correct explication of the input data; the rest of the theory could remain largely unchanged.

Now, before I attempt to characterize the input data in more detail, I want to look at a couple of examples: two cases from the history of science in which theorizing was done in response to the collection of certain data. Here I will be paying particular attention to the actual data that was gathered and the instruments that were used to gather it. The purpose of presenting these examples is twofold. First, I hope that the examples themselves may provide a better sense of what I mean by the “input data” to scientific theorizing. Second, on the basis of these examples I will try to extract a more general characterization of the input data.

4.1 The Geiger-Marsden Experiment

In the period between 1908 and 1913, Hans Geiger and Ernest Marsden performed a series of experiments under the guidance of Ernest Rutherford.²⁶ Their goal was to investigate to what extent rays of α particles could be scattered by various types of matter. To detect the α particles, they used a zinc sulfide screen, which emits light when it is struck by ionizing radiation. Their earliest experiment involved comparing the scintillation patterns on the screen when α particles were shot toward it in a tube filled with air, and then in an evacuated tube. Geiger noticed that the scintillation pattern was more diffuse when air was present in the tube, and Rutherford suggested that this was because air molecules were slightly scattering the α particles en route to the screen.

²⁶Rutherford himself did not have the patience for data gathering. Describing one of his early scattering experiments with Marsden, he wrote: “Scattering is the devil[...] Geiger is a demon at the work of counting scintillations and could count at intervals for a whole night without disturbing his equanimity. I damned vigorously and retired after two minutes.” (Eve, 1939, p. 180)

Several subsequent experiments involved firing α particles at a sheet of metal foil. The first of these showed that α particles were occasionally deflected at angles greater than 90° when they struck the foil. This fact greatly surprised Rutherford. The predominant view of the atom at the time was due to Lord Kelvin and J. J. Thomson. Often called the “plum pudding model,” it maintained that the atom was a positively charged sphere which contained smaller electrons distributed throughout it. According to Rutherford’s calculations, the plum pudding model was unable to explain why such large deflections occasionally occurred. Even on the assumption that an α particle was repeatedly deflected in the same direction by numerous atoms, the chance of a deflection greater than 90° was vanishingly small. Thus Rutherford began to suspect that the occasional large deflections were due to α particles colliding with an intense, localized charge at the center of the atom. Meanwhile, he hypothesized that smaller deflections, which were much more frequent, occurred when the α particles passed through a comparatively large, diffuse charge surrounding the atom’s center.




Figure 1: The Geiger-Marsden Apparatus

On the basis of this hypothesis, Rutherford derived a quantitative prediction about how the frequency of scatterings should vary with the angle of deflection.²⁷ This prediction was verified by experiment in 1913. The setup of the experimental apparatus is diagrammed in Figure 1 above.

Here, R is the source of the α particles, and D is the opening through which they are emitted. F is the sheet of metal foil, for which they used both gold and silver. Together, F , D , and R are anchored in place by the tube T ; they cannot rotate. The zinc sulfide screen is S , which is affixed to a microscope M , allowing for detailed examination of the scintillations on the screen (indicating α particle hits). The microscope is mounted in a box B , which can rotate about the central vertical axis along T . This allows the experimenters to watch the screen for scintillations at varying angles.

Geiger and Marsden of course expected to observe the most scintillations when S , F , and R all lay on a straight line. This is what they observed, indicating that most of the α particles were passing right through the foil without being appreciably deflected. They also expected the number of scintillations on the screen to decrease as the angle formed by S , F , and R became smaller, i.e. as they rotated the screen around the central axis toward the source of the α particles at R . And again, this is what they observed, indicating that fewer α particles were deflected from the foil at extreme angles. However, even when the angle formed by S , F , and R was very small, they still detected a number of scintillations on the screen—far more than would be expected on the plum pudding model, and roughly equal to what Rutherford's model predicted. Now, one might wonder whether these scintillations observed at extreme angles were perhaps due to α

²⁷Rutherford derived several other predictions as well, including the manner in which scattering frequency should depend on the thickness of the foil and on the magnitude of the central charge of the atoms in the foil. For simplicity I will only focus on the relationship between scattering angle and frequency here.

particles leaking through the encasement around R instead of passing through the opening D . And in fact in their original paper, Geiger and Marsden noted that even when there was no metal foil present, they still detected a few scintillations when the angle formed by S , F , and R was small. They explained this by supposing that a few α particles were indeed leaking through the lead casing around R . However, putting the foil back in place at F resulted in a significant increase of scintillations on the screen. Their calculations were thus done relative to the control experiment in which the foil was absent. Even accounting for the low background noise, the number of scintillations observed at extreme angles easily favored Rutherford's model over the plum pudding model. This experiment was thus taken to establish that the atom fit Rutherford's model: it contained a small positive charge in the center, surrounded by a diffuse negative charge.²⁸

Here we have a theoretical result about atomic structure confirmed by a rather straightforward experiment. Our question is: what are the input data? In other words, what are the data that confirmed Rutherford's hypothesis about atomic structure? Well, what are the data that they *reported*? The first data table included in their 1913 paper is labeled "Variation of Scattering with Angle," and it lists the scintillations per minute detected on the screen at various angles, both with and without the presence of the foil. Thus the input data are, most immediately, the scintillations per minute on the zinc sulfide screen at various angles. *These* are the data that their theorizing about atomic structure had to answer to, that it was the purpose of such theorizing to explain. Of course, it is also relevant how the experimental apparatus was set up, how the scintillations were counted, and so on. The important point, however, is that Rutherford, Geiger, and Marsden were not simply handed the structure of the atom, as if from God on a stone

²⁸Actually Rutherford did not know initially whether the central charge was negative or positive. Later experiments confirmed that it was positive.

tablet. They inferred it, and confirmed it, using the results of these experiments. If they were “handed” anything, it was these results (though I’m sure they would recoil at the suggestion these results were “handed” to them). These results are what constitute the relevant input data.

4.2 The Discovery of the Omega Minus Baryon

The second experiment I want to discuss is the discovery of the Ω^- particle.²⁹ In 1961, physicists Murray Gell-Mann and Yuval Ne’eman independently proposed a group-theoretic classification scheme for hadrons (a class of particles that were then thought to be elementary). Their scheme became known as the “Eightfold Way,” and it predicted the existence of a hitherto undetected particle dubbed the Ω^- baryon. Within three years, experimenters at the Brookhaven National Laboratory had announced the discovery of the Ω^- as the result of investigations performed using the Alternating Gradient Synchrotron (AGS) particle accelerator.

In what follows, I want to review the first bubble chamber photograph that constituted the discovery of the Ω^- baryon. But first some background. Particle accelerators work by using a series of magnets to accelerate a beam of particles to nearly the speed of light. These particles are then directed toward a “target” material (which is sometimes a counter-rotating particle beam), and various sorts of detectors are used to measure the trajectories of the particles that result from the collisions of the beam with the target. In the case of the Brookhaven experiment, the accelerated beam was composed of protons, and it was directed toward a stationary target made of tungsten. Upon impacting the tungsten, the proton beam produced a variety of different particle types. These secondary particles were selectively screened until only one type remained:

²⁹See Fowler and Samios (1964) for an excellent review of this experiment.

the K^- meson. The stream of K^- mesons was then directed toward a second target, this time a bubble chamber containing liquid hydrogen.

A bubble chamber works by keeping a liquid heated to just below its boiling point. When particles enter the chamber, the pressure is decreased, at which point the liquid essentially “wants” to boil. Ionizing radiation passing through the chamber provides the nucleation points for bubble formation, so if one takes a picture of the inside of the chamber immediately after the pressure is decreased and the incoming particles have arrived, one can observe the tracks of any charged particles that have just passed through the chamber. Such an image can be seen in Figure 2 below.

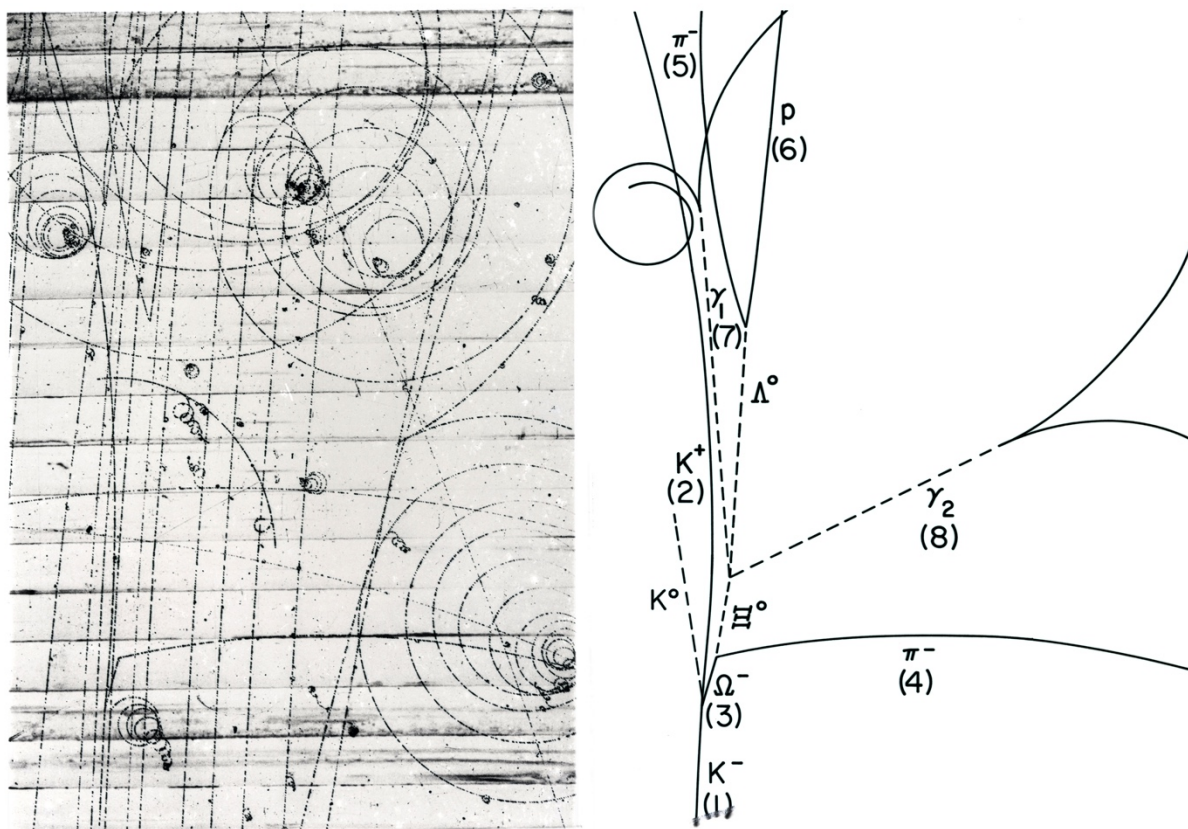


Figure 2: Discovery of the Omega Minus particle. Left: actual bubble chamber image. Right: reconstructed image depicting the relevant event vertices. Courtesy Brookhaven National Laboratory.

I want to focus first on the left image. The tracks here are caused by charged particles passing through the liquid hydrogen in the chamber (neutral particles leave no tracks). The series of nearly vertical lines is the stream of K^- mesons directed toward the chamber. The chamber is subjected to an electromagnetic field, so charged particles moving through it experience a Lorentz force perpendicular to their direction of motion. This results in the spiral tracks visible throughout the picture: negative particles curl clockwise, positive particles counterclockwise.³⁰ Experimenters know the strength of the applied electromagnetic field, so by analyzing the curvature, angles, and thicknesses of the tracks, they are able to identify the particles that produced them.

The Ω^- baryon was identified primarily by its decay products. Calculations suggested that there were three possible processes by which the Ω^- might decay. The one we need to be concerned with here involves the Ω^- decaying into a Ξ^0 baryon and a π^- meson (symbolized by the reaction formula $\Omega^- \rightarrow \Xi^0 + \pi^-$). The Ω^- was also expected to have a very short lifespan, on the order of 10^{-10} seconds.

The right image shows a reconstructed view of the relevant event vertices from the left image. The Ω^- was identified by tracing backward from the observed products. Tracks (7) and (8) are dotted lines, indicating that they are neutral particles not observable in the bubble chamber image. Their orientations can be inferred, and their particles identified, by analyzing their decay products. Notice that each of them decays into a pair of particles with opposite charge—an electron-positron pair. Thus tracks (7) and (8) can be identified as neutral photons, labeled γ_1 and γ_2 . Tracks (5) and (6) were identified as a π^- meson and a proton, respectively (again notice that they have opposite charge). The neutral particle from which they resulted must

³⁰The K^- particles have too much momentum to be appreciably deflected in this short window.

have been a Λ^0 particle, and tracing it back to its origin leads to an intersection with the two inferred photons, γ_1 and γ_2 . At this point, then, some other particle must have decayed into the Λ^0 and the pair of photons. Calculations suggest that this other particle was a Ξ^0 . Now when we infer *its* path, we find that it forms a vertex with track (4), which curves to the right—another π^- meson. Thus the particle on track (3) decayed into a Ξ^0 and π^- : exactly the decay process predicted for the Ω^- baryon. Track (3) is also very short, indicating that the particle that produced it has a very short lifespan, as expected for the Ω^- .

Further confirmation that track (3) was produced by the Ω^- can be gained by analyzing the tracks preceding it. It originated at a vertex with another particle, which produced track (2). This particle appears to have been a K^+ meson. The incoming particle on track (1) was a K^- meson, which was part of the stream of negative Kaons directed at the bubble chamber. Recall that the bubble chamber was filled with liquid hydrogen, so it was essentially a sea of protons (p). The expected reaction for the formation of the Ω^- was as follows: $K^- + p \rightarrow \Omega^- + K^+ + K^0$. So an incoming K^- meson on track (1) must have impacted a proton at the vertex where the Ω^- appears. Thus, out of the expected products we have accounted for the Ω^- and the K^+ . The K^0 , being a neutral particle, did not show up in the bubble chamber, but its inferred path is shown on the diagram, branching off to the left from the same vertex as the Ω^- and K^+ .

So we have a theoretical result: the discovery of the Ω^- baryon and confirmation of the Eightfold Way classification scheme for hadrons. And our question is: What are the input data? And again, the way to answer this is to ask: what are the data that were *reported*? In their original paper published in 1964, a team of researchers at Brookhaven referred to the pair of

images in Figure 2 as their primary source of data.³¹ So the identification of the input data here is straightforward: *they are the tracks observed in the bubble chamber*. Those tracks are what our theorizing has to account for. If tracks like those had never been observed, even after extensive searching, we would regard the Eightfold Way as requiring modification, or perhaps we would even regard it as disconfirmed.

So far, then, we have seen two very different sorts of input data. In the Geiger-Marsden experiment, the input data consisted of the scintillations per minute observed on the zinc sulfide screen when it was positioned at various angles around the apparatus. In the experiments at the Brookhaven AGS, the data consisted of tracks observed in a bubble chamber. I chose these particular examples because of the familiarity of the Geiger-Marsden experiment, and the easily-depictable results of the Brookhaven experiments. Of course, one could multiply examples ad infinitum, but this is not my purpose here (entertaining as that would be). My purpose, again, is to try to convey a sense for what the input data are, i.e. to explicate what sort of data it is the task of scientific theorizing to systematize and explain.

What sort of analysis of the input data is suggested by these examples? It is clear that the input data are comprised of, broadly speaking, observable, measurable phenomena. The scintillations on the zinc sulfide screen can be observed with the naked eye in the dark. And the tracks in the bubble chamber are observable in photographs, and would be easy enough to see if they were not surrounded by the metal walls of the chamber. Of course I am not claiming that *all* input data have to be observable with the naked eye. But they better be observable *somehow*.

But there is an obvious problem in suggesting that the input data are the observable phenomena. After all, I am using the input data as part of a theory of laws: on the current

³¹They also included numerous calculations that were used to identify which particles produced which tracks.

suggestion, the laws are the results of applying the Nomic Formula to the Humean base, and the Humean base consists of the observable phenomena. But the laws themselves evidently determine what does and does not count as observable. For example, if the laws about the propagation of light rays were different, then different sorts of things would be observable, simply because observing something requires light to bounce off of it and reach our eyes. To put the worry sharply, what makes certain facts observable is partly determined by what the laws of nature are, so the laws of nature cannot in turn be determined by what the observable facts are.

To avoid this circularity, we need to find a characterizing property of the input data such that it is implausible that the laws themselves determine the extension of that property. One such proposal is that the input data are *macroscopic phenomena*. In the Geiger-Marsden experiment, for example, the input data were macroscopic scintillations on the zinc sulfide screen. (Of course, the posited *explanation* of those phenomena appeals to microscopic entities, but the phenomena being explained were themselves macroscopic.) Similarly, in the bubble chamber photographs from the Brookhaven AGS, the data consisted of macroscopic bubble tracks.

Now, whereas it is plausible that the laws play a role in determining what is observable, it is less plausible that the laws play a role in determining what is macroscopic. To see this, compare the following two questions:

1. “Why is that observable?”
2. “Why is that macroscopic?”

In many cases, question (1) seems both well-formed and susceptible to an answer that appeals to the laws of nature. For example, one might ask question (1) about a distant galaxy, and the answer would involve appealing to the laws of optics and the construction of the telescope that we used to observe the galaxy in question. One might also ask question (1) about a more

mundane object, like the eraser on the end of my pencil. And here again the answer would plausibly appeal to the laws of optics. By contrast, it is difficult to imagine a case in which question (2) seems well-formed, and it is even harder to imagine a case where it could be answered by appeal to the laws of nature. The most natural case I can imagine where question (2) would be appropriate is one in which an object that is typically microscopic actually turns out to be macroscopic. Suppose, for example, that I encounter a bacterium the size of my pointer finger; then I might reasonably ask question (2)—and with a certain degree of alarm. But whatever the explanation is of the macroscopic-ness of this bacterium, it seems unlikely that it will have to involve appealing to the laws of nature.

All of this is to suggest that I do not think we will have the same circularity worry with “macroscopic” that we did with “observable.” Whereas the latter notion cannot legitimately be taken as part of the *analysans* in a theory of laws, the former notion can. This, then, is my proposal: the input data—the Humean base—consists of the macroscopic phenomena. Consequently, on my view the laws of nature are the results of applying the Nomic Formula to the totality of macroscopic phenomena.

I do not have an explicit analysis, replete with necessary and sufficient conditions, for the concept of macroscopicness. But one way to get a better grip on the concept is to imagine trying to explain a scientific experiment to a competent English speaker who does not already understand the experiment. Imagine, for example, that you are teaching someone how to understand the results of the Geiger-Marsden experiment, yet this person doesn’t know anything about α particles or atoms or charges. How would you teach them to identify the α particles? Probably you would tell them to look at the scintillations on the screen. Since they are a competent English speaker, they know what “scintillations” are: they can detect them, count

them, etc. Similarly, consider how you would teach someone to understand a bubble chamber photograph. In particular, consider how you would teach them to identify the particle trajectories and to figure out their charges. Presumably, you would tell them to look at the tracks to find the particle trajectories, and to analyze the curvature of the tracks to identify the charges. Again, since they are a competent speaker, they know how to identify what counts as the “tracks” in the chamber, and they can easily determine their “curvatures.” In these examples, the data you would tell that person to look for—the scintillations, the tracks, and their curvatures—constitute the macroscopic phenomena.

I fully grant that the foregoing characterization of “macroscopic phenomena” is vague, and that it would need precisification before it can be considered complete. But I actually think this vagueness is a *virtue* of my account. For my purposes, it does not matter very much *exactly* what counts as macroscopic, for the general flavor of my theory will not be greatly affected by precisely where we draw the line. (My view’s main attractions, as it were, lie elsewhere.) Thus the flexibility engendered by the vagueness of “macroscopic” allows readers of different persuasions to precisify it in different ways.

Now, there are a number of potential objections to this account of the Humean base. I will address some of the more pressing ones here, and some other, more well-trod ones in the Appendix.

The most obvious objection is simply that this account of the Humean base is unacceptably anthropocentric. The only reason to privilege the macroscopic phenomena, seemingly, is because they are more readily apparent to creatures like us. But relevance to creatures like us cannot be a good criterion in developing a theory of laws of nature. The laws themselves don’t care about us, so in developing an account of laws, neither should we.

I agree that the account is quite anthropocentric. But I do not see this as a problem. Unfortunately I must hold off on addressing this objection for now. We will have occasion to return to it in the next chapter, where I will consider it in more detail in connection with my pragmatic construal of the Nomic Formula. There I will argue that there are very good reasons for a reductive account of laws such as mine to incorporate anthropocentric elements. For now, though, I have to plead the reader's patience.

Another source of objections against my characterization of the Humean base comes from Cohen and Callender (2009). They canvas a range of possible characterizations of the Humean base, and one of them—the “macrovariable MRL” (for “Mill-Ramsey-Lewis”)—corresponds roughly to the view I am recommending here.

Cohen and Callender present three objections to the macrovariable MRL, which are designed to motivate their own conclusion. Ultimately, they want to argue that we should not “stipulate” once and for all the sort of facts that figure into the Humean base. There are certain advantages to such a move (which they discuss more fully in their paper), though ultimately it strikes me as untenable. Fear of saying something inaccurate should not prevent us from trying to say something informative.

Cohen and Callender's first objection is that theories that posit different conceptions of the Humean base still deserve to be called theories *of laws*. Therefore, our concept of law must not determine, once and for all, the sorts of facts that comprise the Humean base.

This strikes me as an odd line of reasoning. For example, note that one could use this argument against *any* theory of laws in the literature. To do so, just find a place where that theory disagrees with some other theory of laws. Now, since these both count as theories *of laws*, our concept of law must not decide between the two. So the true theory must not determine which is

correct. Clearly such an argument would be specious. Indeed, if every time two theories of X disagreed about something, we could justifiably conclude that our concept of X must not decide between them, then there wouldn't really be much point to philosophical debate.

Their second objection is that any choice of a determinate Humean base seems disagreeably a prioristic. As they put it:

Any particular choice of [Humean base] will preclude (what should be) live empirical possibilities. When we choose [a particular sort of fact to comprise the Humean base], we thereby remove [that sort of fact] from the normal back-and-forth of scientific bartering. This result is especially troublesome for [Humeans] because it is part of their aim to avoid placing constraints on the shape of acceptable scientific theory not justified by broadly empirical inquiry. (\textit{ibid.}, pp. 19-20)

So their point is that, by choosing a particular sort of fact to comprise the Humean base, we thereby isolate that sort of fact from future empirical inquiry; such facts cannot be undermined by later empirical investigations.

I partially agree with this argument, but I do not think it rules out taking the Humean base to be comprised of the totality of macroscopic phenomena. If, as they say, whatever figures into the Humean base is removed from the “back-and-forth of scientific bartering,” then we ought to look for facts that *are* removed (as far as possible) from such bartering. In other words, we ought to look for facts that *could not* be undermined by future empirical investigations. What sorts of facts would those be? The sorts of facts that scientific theorizing, qua scientific theorizing, has to *answer to*: facts about the macroscopic phenomena. Such a conception of the Humean base, I suggest, is not objectionably a prioristic. Essentially, it proceeds by asking, “What sorts of facts must a type of theorizing answer to in order to count as *scientific* theorizing?” What I am suggesting is that if a type of theorizing did *not* answer to the macroscopic phenomena, then it could not count as scientific theorizing. Granted, this is an a priori constraint, but the suggestion is that it is essentially an *analytic* one.

Nevertheless, there is a way of pressing Cohen and Callender's worry here even despite these considerations. John Roberts has suggested that there are certain facts about macroscopic phenomena that *do* seem to be undermined by scientific theorizing. For example, our characterizations of macroscopic phenomena may appeal to a notion of absolute simultaneity: such-and-such a macroscopic phenomenon happened at the exact same time as this other one, for example. But of course, special relativity teaches us that there is no such thing as absolute simultaneity, only simultaneity relative to a reference frame. So here we have facts about macroscopic phenomena that *do* seem to be subject to the back-and-forth of scientific theorizing.

I think this objection mandates a slight revision in how I'm characterizing my theory. Specifically, I need to allow that scientific theorizing can change the way we describe the macroscopic phenomena. So here is how I would suggest we think about it. At bottom, there are the macroscopic phenomena. Using these phenomena as input, science proceeds to construct theories that aim to systematize it. Occasionally, to achieve a nice systematization, scientific theorizing will have to revise aspects of our characterization of the macroscopic phenomena. So the appropriate ways of describing the macroscopic phenomena are not completely determined until the end of scientific inquiry. This conclusion bears some superficial resemblance to Cohen and Callender's conclusion that we should not place a priori constraints on the Humean base. But what is important to keep in mind is that the *general kind* of data that scientific theorizing answers to has not changed. Granted, whether we count a particular sentence as an accurate description of macroscopic phenomena may change as our theories develop, but this is far different from saying that scientific theorizing can change the *kind* of data that it is done in response to. We may revise the ways of describing that data, but not the data itself.

Cohen and Callender's third objection is that stipulating the kind of facts that comprise the Humean base will preclude the development of a theory of special science laws. But I see no reason this should be the case. If we want a theory of special science laws in the same spirit as a theory of fundamental physical laws, then the architecture of the BSA allows for two basic sorts of manipulations: changes to the Humean base and changes to the Nomic Formula. The hope would be that, by manipulating one or both of these elements, we can develop accounts of the laws of various special sciences. Nothing I've said here precludes these sorts of developments.

This completes my defense of the modified Humean base. In sum, my proposal is that we think of the Humean base as the totality of macroscopic phenomena. These are the phenomena that it is the goal of scientific theorizing to predict, explain, and systematize. Such theorizing may eventually tell us that we should revise the way we have been describing these phenomena, but this does not change the phenomena themselves (it only changes our description of them), nor does it change the fact that they constitute the data that such theorizing aims to account for. For those who remain skeptical, I address some residual worries about this characterization of the Humean base in the Appendix. For now, though, I want to provide a more general sketch of my view of laws.

5. A Sketch of the View

The overall view I am advocating can be described very generally, albeit simplistically. The thought is this. When scientists theorize about the laws of nature and the fundamentalia, they do so in response to data about macroscopic phenomena. They gather this "input data," and they then theorize about it using certain principles, and the theories they construct are meant to explain and systematize that data. The totality of the macroscopic phenomena constitutes the

Humean base. And the totality of the principles they use to reason about those phenomena constitutes the Nomic Formula. The laws of nature are the outputs of the Nomic Formula when it is applied to the Humean base. So the overall architecture of the view is the same as Lewis's. The pertinent difference, at present, is that the Humean base is no longer conceived as the totality of facts about the fundamentalia.

On that note, I want to consider an aspect of my view that may not yet be apparent, namely, exactly how it regards the status of fundamentalia (and other microscopic phenomena). Superficially, there is nothing unusual about it. Protons, quarks, virtual particles, etc. all really do exist, according to my view. In this sense, then, the view is straightforwardly realist. However, there is more going on behind the scenes, and it requires some attention.

To start, let's recall how the BSA handles claims about chance. Lewis took chance claims to be a threat to his thesis of Humean supervenience because there was nothing in the Humean base that seemed to deserve the name "chance." After all, Lewis's Humean base consisted solely of non-modal properties, whereas chance seems like a paradigmatically modal notion. Lewis, however, invented a clever way around this. Remember that candidate systematizations of the Humean base are supposed to balance the desiderata of simplicity and strength. Lewis now says that in doing so, these systems are allowed to include claims about chance. The aim of including these claims is to achieve significant increases in simplicity without sacrificing much in the way of strength. For example, imagine that we have a (possibly unfair) coin that we flip one thousand times, and consider the sequence of heads and tails that results. The maximally informative claim about this sequence of results would be a description of each individual outcome: HHTHTTTHTTHTT.... Such a claim, however, is not very simple. Now suppose that in this sequence there are roughly 500 heads and 500 tails. In that case, a candidate system is allowed to

replace the long sequence of heads and tails with the claim that this is a fair coin, i.e. that it has an equal *chance* of coming up either heads or tails. This massively simplifies matters without sacrificing a great deal of informativeness. (Of course, we no longer know the order of heads/tails outcomes, but we can do without these specifics for the sake of the corresponding gain in simplicity.) Now, according to Lewis, what makes chance claims true is if they are included in the candidate system that ends up being the *best* system. In other words, a system may include claims about chance in order to achieve a better systematization of the Humean base. And the truthmaker for a chance claim is the fact that it is included in the best systematization of the Humean base.

Now let's turn from chance to the fundamentalia. First, consider a passing remark made by Barry Loewer:

A proposal that I think captures what many have on their minds when they speak of fundamental physical properties is that they are the properties expressed by simple predicates of the true comprehensive fundamental physical theory. The true comprehensive fundamental physical theory is the minimal theory that accounts for changes of the locations and motions of macroscopic spatio-temporal entities and also for changes in properties that account for locations and motions and so on. (1996, p. 110)

Here Loewer is talking about fundamental physical properties, but the remark would apply equally well to fundamental physical entities. The fundamental physical entities are the entities countenanced by the true comprehensive fundamental physical theory. What is important here is Loewer's explication of the latter phrase. The "true comprehensive fundamental physical theory," he says, is essentially—I am paraphrasing a bit—the minimal theory that accounts for the behaviors of macroscopic entities. In other words (on my reading), *what makes it the case* that some theory is the true comprehensive fundamental physical theory is that it accounts for the behaviors of macroscopic entities. And it is the "minimal" theory that does so, which presumably

means that it does not postulate more theoretical structure than is necessary.³² Now on this view, what makes it true that, e.g., quarks exist? Well, they exist if the true comprehensive fundamental physical theory says they exist. And what makes the claims of that theory, such as the claim that quarks exist, *correct*? *The fact that it accounts for the behaviors of macroscopic entities*. So what are the truthmakers for claims about quarks, on this view? Ultimately, they are made true by their inclusion in the theory that best accounts for the behaviors of macroscopic entities.

My theory takes this passing remark by Loewer and runs with it. On my theory, claims about fundamentalia are treated analogously to how Lewis treated claims about chance. Ultimately, claims about quarks are made true, not by facts about such subatomic entities, but by being included in the best systematization of the macroscopic phenomena.

Thus, it may be tempting to regard my theory as anti-realist about the fundamentalia. Ultimately, it may seem, according to my theory there *aren't* really such things as quarks or electrons, or such properties as spin or charge. In this vein, I want to consider how Ned Hall describes a modified BSA, which I take to be very similar in spirit to mine, though I have reservations about the way he articulates it. The view he outlines takes the Humean base to consist solely of particle locations at times—all other fundamental properties are posited to account for these locations. Here is how he puts it:

The [view] we are now considering holds that the *only* fundamental nonmodal facts there are about [the] world are facts about the *locations* of particles at all times. But a candidate system is now allowed to hypothesize, as it were, that particles are also characterized by additional magnitudes [such as mass and charge], and introduce equations connecting the values of these magnitudes to particle positions[...] What would make it the case that there *are* masses and charges is just that there is a candidate system that says so, and that, partly *by* saying so, manages to achieve an optimal balance of simplicity and informativeness [...] So in our particle world, all there really are are facts about positions of particles at all times; but if we pretend that in addition, each particle is characterized

³²Loewer's minimality requirement can thus be seen as somewhat analogous to Lewis's simplicity desideratum.

by an unchanging value of two magnitudes (one with real values, the other with nonnegative real values), then we can write down very simple equations that encapsulate quite a lot of information about the particle motions. The final step is to let the facts about particle motions that make it the case that these equations achieve such an optimal balance of simplicity and informativeness constitute the truth-makers for claims about particle mass and charge, so that those claims can now be understood as literally correct. What results, again, is a philosophical position about mass and charge that is exactly analogous to the position the BSA already takes about objective chance. Mass, charge, and chance are all, in a certain specific sense, *manufactured* magnitudes. (ms, pp. 27-28, italics in original)

To summarize Hall's view (which he does not officially endorse) in my terms: the Humean base consists of the totality of data about particle positions at all times. We feed this data into the Nomic Formula. All claims about other fundamental properties, as well as the laws of nature, are made true by the outputs of the Nomic Formula: if the Nomic Formula says there are masses and charges, then there are masses and charges. If it says the laws are L_1 , L_2 , and L_3 , then that's what the laws are. But these other magnitudes, and the laws, are in a sense manufactured; all there *really* is are particle locations at times.

There is a great deal to like about Hall's view, though I do not want to get into a discussion of its pros and cons right here. What I do want to focus on is the language of "manufactured" quantities and the claim that particle locations are all there "really" is. There is a sense in which this language can be very helpful in characterizing the view. And it can likewise be used to characterize mine, thusly: all there *really* is are the macroscopic phenomena. The fundamentalia are "manufactured" phenomena—claims about them are made true, not because there really are such things, but because applying the Nomic Formula to the totality of macroscopic phenomena generates—"manufactures"—the fundamentalia as outputs.

While this can be a helpful way to put things, I hesitate to rely on it too frequently because I think it can engender confusions. In short, it misleads by suggesting that the semantics for terms about fundamentalia are the same as the semantics for terms about macroscopic

phenomena. For example, in saying that there really are not such things as protons, but in the next breath saying that there really is such-and-such a track in the bubble chamber, we are presupposing too much of an analogy between the phrases “there is a proton” and “there is a track in the bubble chamber.” The correctness of the assertion “there is a track in the bubble chamber” depends, quite simply, on whether there really is a track in the bubble chamber. But it would be misleading to say that the correctness of the assertion “there is a proton” depends on whether there really is a proton—a minuscule, invisible particle that we can see only indirectly. Rather, the correctness of that assertion depends on a host of other factors removed from the present situation: what other phenomena have occurred in other bubble chambers, what other experimental results have been gathered outside of bubble chambers, etc. To put it provocatively, it is *because* the bubble chamber tracks and a variety of other phenomena consistently behave so much as if they were produced by minuscule, invisible particles that it is correct to say that those particles exist. To put it in Best System terminology: it is because the best systematization of the macroscopic phenomena says there are protons that our claims about protons come out true. But this is far from saying that there really are no such things as protons, or that protons are manufactured entities. Rather, the behaviors of these various macroscopic phenomena constitute precisely the criteria scientists require for concluding that there *really are* such things as protons.

It is important to appreciate the motivations for constructing the view this way. Here I want to look at another passage of Hall’s, this one explaining the motivation for the view he was discussing above. I will quote him at length here because most of what he says applies to my view as well.

[The motivation for the “manufactured quantities” approach] derives from the thesis that the primary aim of physics—its first order of business, as it were—is to account for *motions*, or more generally for change of spatial configurations of things over time. Put another way, there is one Fundamental Why-Question for physics: Why are things

located where they are, when they are? In trying to answer this question, physics can of course introduce *new* physical magnitudes—and when it does, new why-questions will inevitably come with them. (So it is no part of the thesis we are considering that physics is *only* concerned with explaining motions; it is just that the other explanatory demands on it are, in a certain sense, derivative on this one.) Now, what exactly is physics doing, when it introduces these new magnitudes?

One answer is that it is simply postulating new magnitudes, end of story. Perhaps these are fundamental, perfectly natural magnitudes; perhaps they are magnitudes that further investigation will show reduce to more fundamental ones. Regardless, it is part of this realist view that the success of a physical theory that postulates some magnitude at explaining motions—even perfect success, down to the most microscopic level—provides no guarantee that the theory is true in what it says about the given magnitude. For example, a Newtonian theory might perfectly capture the motions of particles in our world, but for all that be false: for it might be that the particles of our world merely behave *as if* they have masses and charges, when in fact they uniformly lack these magnitudes.

A second, very different answer goes instrumentalist: “mass” and “charge” are merely introduced as calculational devices, in a way that utterly deprives claims about their distribution of truth-values. Rival claims about the mass or charge of some particle might lead to better or worse predictions about the motions of it and other particles; but while these *predictions* can certainly be evaluated for truth, the rival claims that partly underlie them cannot.

The version of reductionism sketched in this section can be seen as trying to chart a course between these two answers. With realism, it holds that claims about masses and charges have truth-values. But in an instrumentalist spirit, it holds that the truth conditions for these claims concern, in effect, the predictive success of theories that employ them. In this way, it seeks to avoid the highly implausible view—a view that instrumentalism seems inevitably committed to—that a significant fragment of scientific language does not have the function of describing what the world is like, while at the same time denying the possibility of a certain kind of underdetermination—a kind that can easily seem absurd, given the central explanatory aims of physics. (ms. pp 29-30, emphasis Hall’s)

The only thing I would change about this passage is the claim that physics is fundamentally interested in explaining the spatial configurations of things over time. I think that is somewhat too narrow of a construal of the primary aim of physics. When I dip a pH strip into a liquid and see that it turns red rather than violet, the explanatory demand here is not in any obvious way about the *spatial configuration* of anything. What physics (or chemistry) needs to explain is why the pH strip turned one color rather than another. Now, perhaps it *does that* by appealing to the

spatial configurations of the particles constituting the pH strip, but those are what do the explaining here, not what needs explaining.

As I have already argued, the primary explanatory demand of physics is to account for the macroscopic phenomena, which includes not just the spatial configurations of things over time, but also things like the colors and brightnesses of various indicators (such as pH strips or zinc sulfide screens). Alter that remark in the above passage, and I am fully on board. When physics introduces new entities and magnitudes, it does so to account for the macroscopic phenomena. Claims about these entities and magnitudes are made true if they figure into the best systematization of the macroscopic phenomena. As Hall says, this position avoids the severe underdetermination worries that plague hardcore realist positions (and which, after all, rarely ever exercise *scientists*), yet unlike instrumentalism, it also allows claims about microscopic entities and properties to have truth values.³³

There is another motivation for this view that Hall does not explicitly address. This view essentially aligns our criteria for *figuring out* facts about the fundamentalia with the truth conditions for claims about fundamentalia. In other words, it performs the “epistemic-to-constitutive elevation” with respect to not just the laws, but also the fundamentalia.

Why is that a good thing? Why *should* a theory align our epistemic criteria about fundamentalia with the truth conditions for claims about fundamentalia? Many of the reasons it was attractive in the case of laws apply here as well. Most significantly, it makes sense of why

³³It is worth belaboring the point about underdetermination a bit more. What my theory rules out is the sort of skeptical possibility (articulated, for example, in Smart (1968, p. 150)), in which the macroscopic phenomena all behave *as if* there are microscopic particles with properties like mass, charge, spin, etc., but in reality there actually are no such things, or (perhaps) there are such things but they are vastly different from what is suggested by the totality of macroscopic phenomena. Smart argued that this would be a massively implausible coincidence. According to my view it is not just implausible, but impossible.

scientists appeal to considerations like simplicity in their investigations into the fundamentalia.³⁴

On a more realist picture, it would be unclear why scientists should appeal to simplicity as a guide to the truths about fundamentalia. For why should we expect the fundamentalia to be simple? Again, no reasons are forthcoming; we are back to the unenviable task of giving an epistemic justification of simplicity. But if the facts about the fundamentalia are partly *determined* by considerations of simplicity, this problem dissolves. We again get off easy where other theories encounter an intractable problem.

One more issue deserves comment here. Since the Nomic Formula is now conceived as determining, not just the laws, but also the fundamentalia, it should also be conceived as indirectly determining all the particular matters of fact about the fundamentalia. In other words, the Nomic Formula is not just the truthmaker for claims like “quarks exist.” It is also the truthmaker for claims like “There are four quark pairs located in spacetime region *l*.” All of the particular truths about the spacetime distribution of fundamentalia are made true in virtue of being consequences of a physical theory that is best able to systematize all the particular facts about macroscopic phenomena.

I want to return now to the task with which I began this section: giving a more general overview of my theory. One way to get a feel for a theory of *X* is to ask what sorts of errors about *X* the theory countenances. In the case of laws, the question becomes: What sorts of errors does the theory say are possible for scientists to make when they are investigating the laws of nature? My theory countenances two sorts of errors. The first sort, which we might call “data errors,” correspond to cases in which scientists gather more input data and find it to disconfirm or refute their former theories. The second sort, which we might call “reasoning errors,”

³⁴And they manifestly *do* appeal to such considerations. See the above quote from Gell-Mann in footnote 16, for example.

correspond to cases in which scientists somehow improve their application of the Nomic Formula, and find it to dictate the choice of a different theory. There is no third type of error or mistake we can make about the laws, on this view. So if we imagine (per impossibile) having *all* of the input data, and applying the Nomic Formula *perfectly*, the outputs of that process are, indisputably, the laws of nature.

In general, the fewer potential kinds of error a theory countenances, the more reductive it is. A hardcore nonreductivist theory of laws would allow for at least one more type of error, namely the error in which we have all the input data and apply the Nomic Formula perfectly, but *still* get it wrong. My theory eschews this third type of error, but it explains why it may be tempting to think that such an error is a possibility. It is not possible for us limited beings to ever gather all of the input data, so we will never be in a position in which we can definitively announce that we will never encounter another data error. Thus we can never truly guarantee that we have correctly identified the laws. This fact can make it seem like the laws “stand behind” the macroscopic phenomena. Thus we may be tempted to think that even if we knew *all* of the macroscopic facts, we *still* couldn’t be sure that we had correctly identified the laws. This possibility, I maintain, is merely illusory.

6. Conclusion

By way of wrapping up this chapter, there is one final question I need to address: does my reconceptualization of the Humean base deserve its name? Is it *Humean*? Granted, there is a sense in which the answer does not matter too much. It is just a name, after all, and I have neither the expertise nor the space to get into an extended Hume exegesis. Thus it would be inappropriate to make such a big deal out of the issue. But it is at least worth asking what was

Humean about Lewis's original BSA. The BSA was part of Lewis's overall metaphysical program, according to which all the truths about worlds like ours supervene on the mosaic of point-size quality instantiations throughout all of spacetime. Or as Lewis sometimes put it: "all there is to the world is a vast mosaic of particular matters of fact" (1986, ix). The Humean-ness of the view came from the denial of any necessary connections among these qualities, which were supposed to be entirely loose and separate: "just one little thing and then another."

Are the macroscopic phenomena likewise loose and separate? Are there necessary connections among such phenomena? I am not entirely sure how to answer this question, because I am not entirely sure that I understand the notion of a necessary connection. Certainly there are *constant conjunctions* among the macroscopic phenomena. There are regularities, for example, among the patterns of tracks observed in bubble chambers. Every time we observe a curved track emerging from a vertex with no preceding track, we also observe another track emerging from that same vertex with opposite curvature. (This is guaranteed by the conservation of charge.) But does this count as a necessary connection?

One of Hume's points in Section 7 of the *Enquiry* is that we never *observe* necessary connections, for there are no impressions or experiences from which the idea of a necessary connection can be copied:

In reality, there is no part of matter, that does ever, by its sensible qualities, discover any power or energy, or give us ground to imagine, that it could produce any thing, or be followed by any other object, which we could denominate its effect. Solidity, extension, motion; these qualities are all compleat in themselves, and never point out any other event which may result from them. The scenes of the universe are continually shifting, and one object follows another in an uninterrupted succession; but the power or force, which actuates the whole machine, is entirely concealed from us, and never discovers itself in any of the sensible qualities of body. We know that, in fact, heat is a constant attendant of flame; but what is the connexion between them, we have no room so much as to conjecture or imagine. It is impossible, therefore, that the idea of power can be derived from the contemplation of bodies, in single instances of their operation; because no bodies ever discover any power, which can be the original of this idea[...] therefore

external objects, as they appear to the senses, give us no idea of power or necessary connexion. (EHU 7.8-9: SBN 63-64).

Now, my own understanding of necessary connections is embarrassingly lacking, so if Hume says they are not discoverable by observation of the behaviors of external objects, then far be it from me to contest him. And of course, I can see no reason why there would be any difference for the *macroscopic* phenomena, since they roughly correspond to the sorts of things that we can observe. Thus, it seems safe to say that there are no necessary connections among the macroscopic phenomena. The totality of such phenomena still deserves to be called the “Humean base.” No need for a rebranding.

This concludes my discussion of the Humean base. In the next chapter I examine the Nomic Formula in more detail. I argue that Lewis’s account of it requires significant modifications and additions. My aim is to explicate the desiderata that figure into the Nomic Formula, and to offer a rationale for why scientists’ epistemic principles are constituted by such desiderata.

CHAPTER 3: THE NOMIC FORMULA

1. Introduction

In the preceding chapter, I developed a revised conception of the Humean base. According to my view, the Humean base consists of the totality of macroscopic phenomena, and the laws of nature are the results of applying the Nomic Formula to the Humean base.

This chapter deals with the Nomic Formula, and my main task is to explicate the standards that figure into it. Of course, Lewis has already tried to do that. According to Lewis, the Nomic Formula aims to systematize the Humean base by balancing the desiderata of simplicity and strength. But remember that for Lewis, the Humean base consists, not of the macroscopic phenomena, but of the totality of facts about fundamentalia. So according to him, the Nomic Formula is designed to systematize the totality of facts about the fundamentalia, and candidate systematizations thereof are judged by how well they jointly achieve the desiderata of simplicity and strength. Conversely on *my* view, the Nomic Formula is designed to systematize the totality of facts about *macroscopic phenomena*. And the primary goal of this chapter is to explicate how candidate systematizations of the macroscopic phenomena are judged. One might expect that these differences in our conceptions of the Humean base would engender corresponding differences in our conceptions of the Nomic Formula, and indeed that is the case. On my view the Nomic Formula does not involve (or does not *primarily* involve) a balance of simplicity and strength.

While the implementation is vastly different, some overall features of my view are the same as Lewis's. The Nomic Formula is still conceived as the collection of principles that scientists use to figure out the laws of nature.³⁵ But at the same time, the Nomic Formula is not merely an epistemic guide to the laws. It is not as if the Humean base is the evidence, and the outputs of the formula's operation on that evidence are the things that *it would be most rational to believe* are the laws of nature. Rather, the outputs of the Nomic Formula *are* the laws of nature, and they are so *because* they are the outputs of the Nomic Formula. The laws are whatever the Nomic Formula says they are. So the "epistemic-to-constitutive elevation," which was integral to Lewis's account, is still very much a part of my view.

This chapter proceeds as follows. In §2 I discuss a sometimes-useful way of thinking about the Nomic Formula in terms of Ned Hall's concept of a "limited oracular perfect physicist," or "LOPP." In §3, I introduce a novel problem for Lewis's explication of the Nomic Formula (i.e. the balance of simplicity and strength). On the basis of this problem, in §4 I develop an alternative explication of the Nomic Formula that (i) avoids the problem of §3, and (ii) incorporates desiderata aimed to maximize the laws' *predictive utility*. In §5, I use the results of §4 to develop an argument against alternative theories of laws, and I consider some objections. I conclude in §6 by summarizing some of the primary advantages of my account of the Nomic Formula.

³⁵The phrase "principles scientists use to figure out the laws of nature" might engender confusion. Scientists of course use a wide variety of principles in designing experiments and gathering data about macroscopic phenomena. Those are not the principles I mean to refer to when I speak of the Nomic Formula. Rather, the Nomic Formula consists of the principles scientists apply *to* that data in order to figure out what the laws are.

2. The LOPP

In his manuscript, “Humean Reductionism about Laws of Nature,” Ned Hall distinguishes two guiding ideas behind the Humean account of laws:

The official idea is that the laws are, or correspond to, true statements that collectively encode, in a highly efficient manner, a large amount of information about the world. The unofficial position is a kind of “ideal observer” view, according to which the fundamental laws are whatever a suitably placed observer, implementing the best scientific standards for judging what laws are, would take them to be. (ms, p. 11)

The literature on laws has traditionally focused on the official idea—just note how commonplace discussions are about the balancing of simplicity and strength (the goal of which is to achieve highly efficient summaries of the Humean base). But relatively little attention has been paid to the unofficial idea. Consequently, a good deal of energy has been spent attacking and defending aspects of the official idea that would have been viewed as less important if we had been taking the unofficial idea more seriously. This is unfortunate, for the unofficial idea is significantly more flexible, and also significantly more *promising*, than the official idea.

How is the unofficial idea more flexible? The unofficial idea makes reference only to “the best scientific standards for judging what the laws are,” and it does so without specifying the content of those standards. The official idea, on the other hand, can be viewed as a specific proposal for what those standards are: they involve a balance of simplicity and strength. Now there may very well be problems with the official idea—and we will see a number of them in the next section—even if the unofficial idea is perfectly correct. For the official idea, as it were, sticks its neck out much more: it proposes definitive standards that scientists use for judging what the laws are, whereas the unofficial idea is committed only to the existence of *some* standards or other.

Hall provides a useful heuristic for thinking about the unofficial idea. Here is how he articulates it:

[L]et us imagine someone whom I will call a Limited Oracular Perfect Physicist. What makes our LOPP a perfect physicist is that, given as evidence any information about the world, she is perfectly able to judge what hypotheses about the fundamental physical laws are most strongly supported by that evidence. What makes her oracular is that she has, as evidence, quite a lot of information about the world. But not, of course, all information: else her job would be too easy. (For example, she doesn't directly receive as evidence information about what the fundamental laws are.) Specifically, we will suppose that what she has available to her as evidence are all the facts about the [Humean base] [...] The second guiding idea, then, is roughly that the laws are whatever she says they are. (*ibid*, p. 15)

The LOPP can thus serve as a sort of placeholder for a full-fledged account of the Nomic Formula. Having the idea of the LOPP in mind makes it clear why challenges to the particular formula advanced by Lewis rarely seem to exercise supporters of the BSA very much, and why, as Woodward puts it, supporters of the BSA can seem to think that “something like the BSA ‘must’ be correct” (2014, p. 104). Something like the BSA must be correct in at least this limited sense: there must be *some* standards that an ideal physicist would use to generate a theory about what the laws of nature are. That much should be uncontroversial. The unofficial idea simply adds that the theory of laws that best satisfies those standards *thereby* constitutes the truth about the laws of nature.³⁶

The LOPP can also serve as a useful tool at several points in the dialectic. For example, consider part of my argument against Lewis's construal of the Humean base and in favor of my own. I argued that physicists never apply the Nomic Formula to facts about fundamentalia; those

³⁶What if a radically subjective form of Bayesianism is correct, such that there *are no* particular standards one has to implement when reasoning about the laws; all that is required is that one's beliefs about the laws are probabilistically coherent. (Thanks to Marc Lange for pressing this objection.)

I find such a view quite implausible, at least with respect to the laws of nature. In §4, I am going to articulate a number of the standards that I believe figure into scientific reasoning about the laws, and I will suggest that these standards are reflected, to a remarkable extent, in the actual laws that are posited in scientific practice. I will also suggest a rationale for why *these* standards are the ones that we use. All of this is very hard to make sense of on a view that says there are only minimal probabilistic constraints on our beliefs about the laws.

facts are part of the *outputs* of the formula. So if we are to imagine the LOPP as an ideal physicist, then it is puzzling why she should require information about the fundamentalia in order to tell us what the laws are. Rather, she should rely on the same kind of information that actual physicists rely on: facts about macroscopic phenomena. What makes her “oracular” is just that she has a lot *more* of that information than ordinary physicists do, not a different *kind* of information altogether.

3. Problems with Lewis’s Formula

Let us turn now to examine Lewis’s explication of the standards in the Nomic Formula. As we have already discussed, Lewis maintains that the Nomic Formula involves a balance of simplicity and strength. Lewis does not say much about precisely how these desiderata are to be construed, but at least this much is clear: by balancing the two, we are supposed to come up with a collection of statements that together constitute an efficient summary of the Humean base. These statements might not all express regularities (some may report particular facts), but only the statements that do express regularities count as the laws. Furthermore, the practice of balancing simplicity and strength is supposed to come from scientific practice. It is supposed to be something that scientists themselves do when they are trying to figure out the laws of nature.

This characterization of the Nomic Formula is defective for several reasons. One reason is pointed out by Roberts (2008, pp. 6-10). Roberts imagines a universe that consists just of our solar system, and in this imaginary universe the planetary orbits exactly obey Kepler’s laws. The particular facts to be systematized in this universe, let’s suppose, are facts about the exact positions of all the planets at every moment in the history of the universe. The strongest possible systematization of these facts would just include *all* of them. But that system is not at all simple:

it is a laundry list of facts about planetary positions at times, and it documents no regularities among these facts. Roberts suggests that a much simpler system, but one that still possesses a great deal of strength, would just consist of Kepler's laws. These laws convey a significant amount of information about the history of the universe, and they do so in a much simpler way than the laundry list of *all* the facts. Granted, they are not as informative as that list, because they do not entail any particular planetary positions at particular times. But we are to suppose that this deficiency in strength is more than made up for by the corresponding increase in simplicity.

The problem arises if we compare this second systematization, consisting of just Kepler's laws, with a third systematization, this one containing Kepler's laws as well as some information about planetary positions. In particular, Roberts imagines that we add enough information to Kepler's laws to be able to deduce the positions of all the planets at all times. We could, for example, add equations expressing the particular trajectories of all the planets in the solar system. Disregard, for a minute, any considerations about whether these equations should constitute laws of nature. What Roberts points out is that the comparison between these two systems, one that just consists of Kepler's laws, and the other which appends the particular planetary equations of motion, is going to be problematic. Lewis, of course, directs us to compare them by seeing which comes out higher in the balance of simplicity and strength; and to explicate those standards we are supposed to look to scientific practice. The problem, Roberts suggests, is that comparisons of the requisite sort are *never done* in scientific practice:

In choosing which theories to accept, scientists don't ever face a choice like the choice between the two systems in the toy example. One of those systems is just Kepler's laws; the other is Kepler's laws together with the equations of motions of all the planets. The two are not incompatible with one another; in fact, the second one entails the first one. Scientists might have to face the question of whether they have sufficient evidence to justify accepting the stronger theory, or whether they should be more conservative and merely accept the weaker theory. But this is a judgment about the strength of the available evidence, and not a judgment about the competing theoretical virtues of two

systems. So it is hard to see why we should expect that scientific practice will ever need to rely on standards for balancing information content against simplicity, in the way that Lewis's best-system account of laws requires. (p. 9)

In short, then, the problem is that Lewis's explication of the Nomic Formula directs us to appeal to scientific standards that do not exist. We *have no* standards for comparing the two sorts of systems Roberts was considering. Scientific practice does not have the resources that Lewis's account requires.

A closely related problem, pointed out by both Hall (ms) and Woodward (2014), is that it looks like Lewis's formula—regardless of whether it can be found in scientific practice—will deem certain statements to be laws that intuitively are not. Let's stick with the Keplerian solar system, and the systematization containing solely Kepler's laws. Roberts suggested adding the equations of motion for all the planets, thereby allowing us to deduce all the facts about planetary positions at all times. Now, it's possible that Lewis would have objected that these multiple equations are too complex to justify their inclusion. Of course, one pertinent question is, "too complex" according to *what*? After all, Roberts has already argued that scientific practice cannot underwrite such a verdict. Nevertheless, suppose there is some legitimate basis on which Lewis can rule out inclusion of the equations of motion for the planets. Even still, there are other statements that could likewise be used to nail down the exact history of this Keplerian universe, and these other statements are simple by *any* metric. For example, since Kepler's laws are deterministic in both temporal directions, all we would need to add is a statement about the positions of all the planets at some particular time. We could even wait until some time when the positions of the planets could be stated very simply, e.g. when they all "line up." By including such a statement, a candidate systematization would buy an enormous amount of informativeness at a small cost in simplicity. In a similar vein, Hall (ms, p. 45) suggests the possibility of a

candidate systematization including a “phony fundamental constant,” i.e. a constant that encodes the exact physical state of the universe at some instant in a single real number. The inclusion of such a constant in any systematization with forward- and reverse-deterministic dynamics would nail down the exact history of the entire universe, thus vastly increasing the system’s informativeness without drastically diminishing its simplicity.

Now, the worry is not that the initial conditions (or the phony fundamental constant) *themselves* will wrongly count as a law of nature, according to Lewis’s view. For there is no straightforward sense in which they constitute regularities. Rather, the inclusion of such facts in the best system along with Kepler’s laws will allow us to derive all sorts of *other* regularities that Lewis’s theory will be forced to count as laws, but which manifestly are *not* laws. For example, they will allow us to derive regularities about the orbits of the particular planets, and since these constitute regularities and are included in the best system, by Lewis’s lights they must count as laws. But in actual scientific practice, statements describing particular planetary orbits are *not* regarded as laws. In sum, then, the problem is this. A statement detailing the total physical state of the universe at an instant would constitute a very economical way of conveying information about the rest of that universe’s history (in conjunction with statements of Kepler’s laws). So by those lights, it should be included in the best system. However by including it, we generate a number of theorems about regularities in the behaviors of the planets that should not count as laws, but which the BSA must count as laws. Lewis’s explication of the Nomic Formula thus renders incorrect verdicts about the laws of nature.

I agree with the foregoing objections posed by Roberts, Woodward, and Hall, but I do not think they cut to the heart of the matter. Rather, they are symptomatic of a deeper problem with

Lewis's approach. To see this, it will first help to isolate the following three features of Lewis's account of the Nomic Formula:

1. It is supposed to originate in scientific practice.
2. It involves a balance of simplicity and strength.
3. It is designed to generate efficient summaries of what goes on in the Humean base.

It will be useful to consider the relationships between these features. What, for example, is the relationship between (2) and (3)? Well, (2) is manifestly designed to *achieve* (3). In other words, the point of balancing simplicity and strength is to generate efficient summaries of what goes on in the Humean base. This point has been made by several authors, most entertainingly by Lange, who explicates the motivation behind balancing simplicity and strength as follows:

You: Describe the universe please, Lord.

God: Right now, there's a particle in state Ψ_1 and another particle in state Ψ_2 and I'll get to the other particles in a moment, but in exactly 150 million years and 3 seconds, there will be a particle in state Ψ_3 and...

You (checking watch): Lord, I have an appointment in a few minutes.

God: Alright, I'll describe the universe in the manner that is as brief and informative as it is possible simultaneously to be—by giving you the members of the “Best System.”

You: Do tell... (2009, pp. 101-102)

It is clear here that Lange takes the purpose of balancing simplicity and strength to be achieving an efficient summary of the entire universe. Furthermore, Lange's quote also clarifies how Lewis is understanding the relevant sense of “efficiency.” These are supposed to be summaries that are efficient from a *God's-eye* perspective: essentially, they are the sorts of things that you would look for if you had an academic curiosity about what goes on in the Humean base, and you could see the entire thing. In short, Lewis's laws are “God's-eye-efficient summaries.”

With that in mind, what is the relationship between features (1) and (3) of Lewis's formula? I think the way we should understand this is that Lewis is giving us a story about the goal of scientific theorizing about the laws. He thinks scientists (physicists in particular) theorize

about laws of nature with the goal of achieving God's-eye-efficient summaries of the totality of the particular matters of fact. The issue I want to point out is that it is implausible that this is what physicists are doing, or at least that this is *only* what physicists are doing, when they investigate the laws of nature.

Why is that implausible? Roughly, because physicists are human beings concerned with navigating the world. Given this, there are certain *pragmatic* features that it would be useful for our concept of a law of nature to possess, and which are *very unlikely* to be possessed by a God's-eye-efficient summary. This point is perhaps best illustrated by example.

Suppose you are trying to find your way out of an immense maze, and you come to a point at which you are unsure whether to go left or right. Suddenly, a mysterious stranger appears, and offers to sell you information about the maze to help you get out. The stranger offers you four types of information:

1. Details about the total number of lefts and rights you must take to escape
2. Directions that tell you where to turn based on facts about the *opposite side* of the maze
3. Directions that tell you where to turn based on your distance from the center of the maze
4. Directions that tell you where to turn based on your current surroundings

Each type of information costs the same amount, and you only have enough money to buy one.

Which one do you buy?

Clearly, you buy option (4); it tells you, based on your current surroundings—i.e., information that is readily available to you—where you should go. Conversely, the other options are less helpful. For example, consider option (1). Just knowing the total number of lefts and

rights doesn't tell you what you should do *here*. In fact, it really doesn't tell you what to do at *any* point. Similarly, options (2) and (3) are unhelpful because they rely on information that you, as a person inside this maze, do not have access to.

Note, however, that from a God's-eye perspective there is nothing wrong with options (1) through (3). Imagine, for example, that instead of being inside the maze yourself, you are navigating the maze with a pencil on a piece of paper, so you can see the entire thing at once. If you were in *this* position, then information in the form of options (2) and (3) would be perfectly acceptable. Furthermore, if we suppose that you aren't even trying to navigate the maze, but are merely academically curious about it, then information in the form of option (1) would also be perfectly acceptable. The problem is that, from the perspective of someone *within* the maze, the sort of information expressed in (1) through (3) isn't very helpful. In short, the types of information about this maze that you will find useful depends on your perspective.

My suggestion is that a similar consideration holds about the laws of nature. Physicists don't just want God's-eye-efficient summaries of the entirety of the universe. They want principles that will help them figure out how to *navigate* the universe; they want principles with certain *pragmatic* features. Thus, their epistemic standards are going to reflect this fact, i.e. they will be designed to discover such pragmatic principles. And now remember that, according to the BSA, the Nomic Formula is itself supposed to be an elevation of physicist's epistemic standards to the status of standards that are constitutive of lawhood. So what I am suggesting is that a preference for principles with certain pragmatic features should be built into the Nomic Formula.³⁷

³⁷It might be objected that Lewis's conception of the Nomic Formula is *already* quite pragmatic. After all, an efficient summary is much more practical than a detailed list of every particular matter of fact that obtains. But my point is that what is pragmatic from a God's-eye perspective may not be *at all* pragmatic from *our* perspective.

Now, this talk of “navigating the universe” is all very figurative at the moment, and I will aim to make it much more precise in the next section. But I want to linger for a minute on the motivation for introducing pragmatic elements into the Nomic Formula. To an anti-reductionist, this likely sounds downright outrageous. Why should *the laws of nature* answer to our practical concerns? Only someone like Lewis’s ratbag idealist would even *entertain* such a proposal. However, it is not as outrageous as it may first appear. Here I am again going to draw from Ned Hall’s discussion, for he gives an eloquent articulation of what we might call the “pragmatic turn.” (In the following passage, Hall is discussing a reductionist view according to which the Humean base consists solely of particles and their natural properties. However, nothing substantive hangs on its differences from my view.)

[Let’s] step back for a moment, and reflect on the implications of the basic metaphysical outlook adopted by the reductionist. It helps to keep firmly in mind that he thinks that for any world, *all there is* to that world is a distribution throughout space and time of various perfectly natural magnitudes. For example, all there is to our Newtonian particle world are some particles moving around, with masses and charges. That’s it. It is emphatically *not* that the facts about these particles serve as clues to something “behind the scenes” that is directing their behavior. That is quite the wrong way to think about it. In fact, a much better way to think about the status of laws, given such a background metaphysics, is *pragmatically* [...]

We have an array of nonmodal facts about a world—as it might be, facts about how many particles there are, what their masses and charges are, and how they are moving. Science, presumably, is in the business of investigating *that* stuff. It is not in the business of investigating any further stuff that lies behind the scenes, for the scenery constitutes all of reality. So do not ask what, given such a metaphysical outlook, laws are. Ask instead how, given such a metaphysical outlook, one might usefully draw a distinction between certain facts about the world that are in some sense *distinctively appropriate targets for scientific inquiry*, and other facts that are less interesting and central.

As far as I can see, for a thoroughgoing reductionist about “laws”, the only interesting sense in which that word picks out a subject matter worthy of philosophical attention is if it demarcates a class of claims about the world that are, somehow, distinctively appropriate targets for scientific inquiry. It is, of course, a good and hard question what would make some claim a distinctively appropriate target for scientific inquiry[...] For now, I simply wish to note that there is evidently plenty of room for the

When I say that Lewis’s account of the Nomic Formula is not sufficiently pragmatic, I mean this to be understood from the perspective of *creatures like us*.

view that it is in part facts about us—*idiosyncratically* about us, and our peculiar human psychologies—that play a role in determining the most appropriate way for us to structure our investigation of the world. How could the details of our peculiar human situation *not* be relevant to this matter? (ms, pp. 39-40)

There is a lot to digest here. One of Hall's central points is that given the reductionist's fundamental metaphysical picture of the world—be it a collection of particle motions, or macroscopic phenomena, or whatever—there is a challenge that does not show up for non-reductionist theories of law. The challenge is to explain why laws of nature are, as it were, *worthy of our attention*. According to *nonreductionist* views, this challenge is less apparent. On such views, laws are *automatically* worthy of our attention, for they “control” or “govern” the happenings of the universe. They are, for lack of a better phrase, metaphysically interesting, so *no wonder* we are interested in investigating them. But on reductionist views, it is not immediately clear why physicists should be interested in investigating the laws of nature. After all, on such a picture it is *not* as if the laws are something “behind the scenes,” for, as Hall puts it, “the scenery constitutes all of reality” (p. 39). In other words, according to the reductionist, laws are metaphysically *uninteresting*. They are just patterns in the phenomena. So on the reductionist view, we need some *other* explanation of the laws' attention-worthiness, something that doesn't appeal to their metaphysical status.

As Hall suggests, one plausible explanation for our interest in the laws compatible with a reductionist metaphysics is that there are *pragmatic benefits* to discovering them. This would naturally explain our interest in the laws—if it's useful to know them, why *wouldn't* we want to know them?—and it does not require them to be metaphysically privileged in any way. Of course, to say that there are pragmatic benefits to discovering the laws is not to say what those benefits are, and any legitimate reductionist account would have to go much further in articulating exactly what is pragmatic about the laws of nature. But the basic point remains: by

attempting to understand the laws pragmatically, the reductionist is able—hopefully—to explain why we are *interested* in discovering the laws of nature. So the appeal to pragmatism is, I think, quite appropriate for the reductionist.

Are there other routes the reductionist might take to account for our interest in the laws? There are, but they do not strike me as anywhere near as plausible as the pragmatic route. Here are two alternative possibilities for the reductionist, along with reasons I do not think they will work.

One potential explanation for our interest in the laws, compatible with a reductionist metaphysics, is that our interest is primarily *aesthetic*. After all, the laws have a characteristic generality, even a sort of *elegance*, so perhaps it could be maintained that these two features together explain our interest in them. However I do not think this is a very compelling explanation, and it would be better if we could find more *forceful* reasons behind our interest in the laws. Surely, for example, it is not just for aesthetic reasons that we spent over \$13 billion investigating the Higgs boson.³⁸ It would be a very tall task to explain that endeavor solely, or even primarily, by reference to the elegance of the resulting theoretical upshots (elegant though they are). Of course this is not to say that aesthetic reasons play no role whatsoever in our interest in the laws. But it is to suggest that they play at best a subsidiary one.

Alternatively, the reductionist could suggest the following explanation for why physicists are interested in discovering the laws of nature: they are confused. Perhaps they think of themselves as investigating God's commandments for the natural world, or perhaps they implicitly adopt something like a non-reductionist conception of lawhood, according to which, again, the laws would *automatically* be worth their attention. No doubt some physicists *do* regard

³⁸“How Much Does It Cost To Find A Higgs Boson?” (Knapp, 2012).

their investigations in just these ways, at least in their more philosophical moments. But explaining the entirety of physical reasoning about the laws in this manner should be a last resort for the reductionist. For one, this puts her in a dialectically tenuous position; nobody wants to be committed to the claim that the majority of research and activity in the history of physics is based on a confusion. But furthermore, it is exceedingly unlikely that a practice based solely on confusions, that provides no practical benefits, would sustain itself in the competition for funding for scientific research. Money for academic research is in high demand, and if no practical benefits were derived from physics, then we would not observe developments such as the rapid proliferation of multi-billion-dollar particle accelerators in the middle of the twentieth century. So on the reductionist picture, even if individual scientists are sometimes confused about what they are doing when they investigate the laws of nature, we should still expect to be able to find practical benefits that explain why our society invests so heavily in those investigations.

The upshot of all this is that reductionists should seek to understand the laws from a pragmatic perspective. That is, they should seek an account of laws that places primary importance on the practical *uses* to which we put them. Now, one way to articulate my objection to Lewis's explication of the Nomic Formula is to observe that an efficient summary of the Humean base from a God's-eye perspective is not sufficiently pragmatic to be worth our attention. Or at least, it is not worth *as much* attention as physicists tend to give the laws of nature, and thus it is implausible that this is what physicists are looking for when they are trying to discover the laws. This is because, for one, such a summary is likely to contain a lot of information that is pretty useless to us—information about what is going on in distant corners of the universe, for example, or information about what will happen long after we are all dead.³⁹

³⁹Why then, you might ask, do we concern ourselves with things like cosmology? For that discipline concerns itself with exactly the sorts of facts that I have been saying are not particularly useful for us to know. I think it is helpful to

But more importantly, it's perfectly conceivable that in generating a God's-eye-efficient summary of the entire Humean base, the sort of information we are interested in may get either *swamped out* or *obfuscated*. For example, it's possible that the comparatively useless information about distant events will swamp out the more useful information about how things behave around here, and in that case we'll be left with nothing useful.⁴⁰ Alternatively, it's possible that the information we care about won't get swamped out, but instead it will be condensed into a form that makes it unusable. Consider again the case of the maze, from which you are trying to escape. As we saw, there are plenty of ways to articulate details about your escape path—e.g., the total number of lefts and rights that you have to take—that would be of no use to you. But if our goal were to write down a God's-eye-efficient summary of your escape path, information about the total number of lefts and rights may indeed be the best way to do so. So the worry is twofold: Lewis's efficient summaries are likely to (1) contain useless information, and (2) *omit* useful information, e.g. they might give us sketches where we need details, or they might give us details, but in a form that is not particularly useful to us.

In the next section, I am going to develop a revised account of the Nomic Formula that takes into account the sorts of considerations I have just been discussing. The general thought is going to be as follows. There are an infinite number of principles providing information about the totality of the particular matters of fact. Many of these principles would be utterly useless to

see disciplines like cosmology as an outgrowth of branches of physics that *do* study domains that are of more practical concern to us. Roughly, once our standards and methods of physical reasoning are shaped by pragmatic concerns arising in practical domains, they can be applied to other domains as well. And as for our *justification* for being concerned with cosmological questions? I see no problem with saying that we do so simply because we are *curious*.

⁴⁰Case in point: it is estimated that dark energy constitutes 68.3% of the universe's total energy, and yet it plays almost no role in how things behave here on Earth. Consequently, an efficient summary of the totality of the universe will presumably, to our inconvenience, spend a lot of its time talking about dark energy, and less of its time describing the behavior of ordinary energy.

creatures like us. A few, though, would be more useful. My suggestion will be that by thinking carefully about what sorts of features make a systematizing principle like this useful for us, we can generate an account of the Nomic Formula that, unlike Lewis's account, both agrees with our intuitions about what the laws should be and accords nicely with the laws we find in actual physical practice.

4. A Pragmatic Account of the Nomic Formula

How should we try to figure out what sorts of pragmatic standards figure into the Nomic Formula? Here we need not proceed from purely a priori considerations. Rather, we can look to actual scientific practice. The question we should ask to start is: What *uses* do we make of the laws of nature?

One answer that comes to mind involves appealing to the various conceptual roles attributed to the laws. The laws, for example, support counterfactuals, figure in explanations, and underwrite inductive inferences. Of the various roles played by the laws, I think the most explicitly pragmatic is their use in connection with inductive inferences, or, to put it more straightforwardly, their use in predictions. It is undeniable that we use the laws to make predictions, and we plan our actions accordingly. Just think of the sorts of uses that we regularly make out of putative laws: we forecast everything from astronomical events to the weather, design things from spacecraft to nuclear reactors, and develop ever more powerful computers. In all this, the role of predictions licensed by the laws of nature is crucial.

Thus what I want to do is to consider, in the abstract, what it takes for some set of principles to be predictively useful for creatures like us. In other words, I want to consider what sorts of features a set of predictively useful principles would likely exhibit. If I am right that

reductionists should view the laws pragmatically, and that the primary pragmatic use of laws is predictive, then by thinking about what it would take for a set of principles to be predictively useful to creatures like us, we should arrive at a number of the features exhibited by actual (putative) laws of nature.

Two clarifications. First, I am not committed to the *completeness* of the following list of features. While I think these features constitute some of the most obvious and most significant ways in which a set of principles can be predictively useful to us, I do not mean to foreclose the possibility that other features could also figure in.

Second, as with Lewis's Nomic Formula, it will not always be possible to achieve principles that exhibit all of these features simultaneously—that will depend on how kind nature is. Thus it would be more appropriate to call them “desiderata” rather than “features”: the Nomic Formula directs us to look for a set of principles that best satisfies these various desiderata. But the important point is that we will sometimes have to make trade-offs. Lewis's idea of balancing these various desiderata is therefore a useful one to keep in mind, only now we can say a bit more about what this “balance” is meant to achieve. The best systematization of the totality of the particular matters of fact is the one that achieves the best balance of the following desiderata, and the “best balance” is the one with the highest *predictive utility*.

4.1 Informative Dynamics

Suppose that we want to make predictions about a particular system *S*. One thing it would be quite useful to have at our disposal is a dynamical principle that tells us how *S* evolves over time. And the more informative this dynamical principle is, the better. If it tells us uninformative sorts of facts about *S*, like “*S* will change over time,” it will not be very useful. Simply knowing that *S*

will change doesn't allow us to make any substantive predictions or plans regarding S ; for that, we need to know *how* S is going to change. Similarly, this dynamical principle will be more informative if it tells us with *certainty* how S will evolve over time, i.e. if it is deterministic. If instead it only gives us probabilities of how S will behave in the future, it becomes correspondingly more difficult for us to make plans regarding S . So our first desideratum is this: for a set of principles to be predictively useful for us, it should include a maximally informative dynamical principle. In the ideal case this means that the principle is deterministic, though in §4.3 and §4.6 we will see how it could fall short of determinism.

Now, this is not to say that for a set of principles to be predictively useful to us, they must all be dynamical principles. Indeed, strictly speaking, *none* of them *has* to be a dynamical principle. All we require is that, jointly, the set of principles *implies* a dynamics for the system of interest S . This accounts for the fact that some actual putative laws of nature are not straightforwardly dynamical (though of course, some are). For example, the ideal gas law is not explicitly dynamical, though we *can* use it to predict the dynamics of a given system. More generally, the crucial point is that the actual putative laws of nature jointly imply a dynamics for various systems. In fact, the core ingredients of a paradigmatic physical theory are (1) a space of possible physical states (a “state space”), and (2) a dynamical principle (or set of principles) that specifies the trajectory of any given system through state space.⁴¹ It is clear from just this simple description of the fundamental ingredients of a physical theory that the laws of nature are going to have to have extensive dynamical implications.⁴²

⁴¹In classical mechanics, the state space of an n -particle system is a $6n$ -dimensional space, and the dynamical principles (which vary depending on which formulation one uses) are deterministic. In quantum mechanics, the state space is a *vector* space (whose dimensionality varies, but is often infinite), and the dynamical principle is the Schrödinger equation, which again is deterministic.

⁴²Here I am in agreement with Hall (ms, p. 1) and Maudlin (2007, p. 13), who both regard the dynamical implications of the laws as a defining characteristic.

4.2 Wide Applicability

Let's now suppose we want to make predictions and plans about a different particular system, S' . One way for our set of principles to allow for that is to include another dynamical principle, this one applying specifically to S' . And in general, for each new system we want to make predictions about, we could include a new dynamical principle that details how that new system evolves over time. The problem, however, is that this is quickly going to get unwieldy. We humans make predictions and plans about a great many different sorts of systems. It would be massively inconvenient if we had to both discover and then apply a different dynamical principle for each one. A smarter way to accommodate the desire to make predictions about a variety of different systems would be to instead include (amongst our set of predictive principles) just a few dynamical principles, but require that those dynamical principles apply to all (or most) of the variety of systems we want to make predictions about. In other words, we gather a collection of systems S , S' , S'' , etc., that we want to make predictions about, and try to extract a general dynamical principle that applies to all of them.

Of course, to accomplish this without sacrificing the informativeness of the dynamical principle, that principle is going to have to be sensitive to the differences amongst these systems. In other words, suppose a dynamical principle D applies to two systems, S and S' . Now, D *could* just provide us with accurate descriptions of the properties common to the evolution of both systems. But in doing so, it will sacrifice some informativeness about the idiosyncrasies of those systems' temporal trajectories. In order for D to be maximally informative about the future states of both S and S' , D should be sensitive to the differences between them. In other words, the

dynamical principles should admit systems having a variety of initial states or conditions, and their outputs should depend on the details of those initial conditions.⁴³

Already, then, we can see that in order for a set of principles to be predictively useful to creatures like us, it should include principles that are maximally informative about the dynamical evolution of the systems of interest while being very *permissive* about the allowed initial conditions of such systems. Several theorists such as Hall (ms), Woodward (2014), and Wigner (1970) have suggested that the laws of nature we find in scientific practice likewise exhibit these features: they are maximally informative in their dynamics but maximally permissive in the allowed initial conditions of the systems they apply to—including, most notably, the universe as a whole. Now what I am suggesting is that this is no accident.⁴⁴ You would *expect* to find *exactly these sorts of principles* in use amongst creatures who, like us, make predictions and plans about a variety of different sorts of systems.

There are further features that predictively useful principles would likely exhibit aside from the two I have just noted (i.e. informative dynamics with a variety of permissible initial conditions). Suppose, again, that we are trying to use dynamical principle *D* to make predictions about system *S*. Here it would be quite helpful if *D* did not require us to have knowledge of facts *outside* of system *S* that we do not typically have, or that it would be prohibitively difficult to ascertain. For if *D* requires knowledge of such facts, we obviously will not be able to use it to make predictions about *S*. This very simple consideration is going to turn out to have a wide

⁴³Here I do not mean “initial” in the sense of the initial conditions of the entire universe. I only mean to refer to the conditions exhibited by the system at the time we want to make predictions about it.

⁴⁴In the literature on laws, the term “accident” is often used in a somewhat technical sense, to denote any fact that is not entailed by the laws of nature. I do not intend this technical reading here.

variety of consequences about the features exhibited by the dynamical principles that we use for prediction.

4.3 Spatial Locality

Suppose that in order to predict how S is going to behave in the next few minutes, D requires us to plug in information about what's happening somewhere on the other side of the universe.

Now, if our concern were to generate efficient summaries of the totality of facts from a God's-eye perspective, there would be no problem with such a requirement. For it may indeed happen that the most informative, efficient patterns among the facts make reference to correlations between very distant events. From a God's-eye perspective, there is nothing objectionable about this requirement. But from *our* perspective, from the perspective of creatures who want to make predictions about S but have no access to facts on the other side of the universe, this requirement is *quite* objectionable indeed. It essentially precludes us from being able to use D to make predictions about S . Imagine, for example, that Newton's second law were not $F = ma$, but instead $F = m^n a$, where n denotes the total number of particles in the universe. If that were the law, then we'd be out of luck, for there is presumably no way we can reliably figure out the value of n . But it is, I'm suggesting, *no accident* that Newton's second law (as well as every other law that has ever been proposed) fails to include reference to n . A principle that makes crucial reference to the total number of particles in the universe is neither the kind of principle that creatures like us would likely be able to discover, nor one that we could use for making predictions.⁴⁵

⁴⁵Of course, Newton's law obviously *couldn't* be $F = m^n a$, for that would generate all sorts of incorrect predictions. What's important here is not the exact manner in which the laws fail to make reference to n , but that none of them make any reference to it at all.

What I am suggesting, in other words, is that there is a spatial locality desideratum that the dynamical principles should satisfy: a dynamical principle D describing the temporal evolution of a system S should appeal only to facts that are local to S (e.g. they may be part of S itself, or they may be nearby S). And this is not just a brute desideratum, the existence of which has no explanation. Such a desideratum has straightforward consequences about our ability to use such principles to make predictions.

As I just stated it, however, the spatial locality desideratum is too strong. For example, consider the law of universal gravitation. It states that the force due to gravity on a particle is dependent on the positions and masses of every other particle in the universe. That constitutes a flagrant violation of the spatial locality desideratum as I just stated it. Take S to be some particle, and suppose we want to figure out the force due to gravity on this particle. According to the law, the answer is going to depend on *many* facts that we could never hope to ascertain. This would seem to suggest that the law of universal gravitation is useless to us. But of course it is not useless, for it has the following convenient feature: the force between any two particles is inversely proportional to the square of the distance between them. Thus as we get further and further away from the particle of interest, i.e. S , the influences of the other particles are going to drastically diminish and *cancel out*, such that we can effectively ignore them. In other words, we can get a very reliable approximation to the correct answer by ignoring all of the facts that we have no way of figuring out.

What this shows us is that we need to restate the spatial locality desideratum on dynamical principles, thusly: a dynamical principle D describing the temporal evolution of a system S should either (1) not require information about facts that obtain far away from S , or (2) accord diminishing influence to facts as they become further and further from S , such that facts

that are too distant to ascertain can be effectively ignored. In short, D should treat S as at least quasi-enclosed. This desideratum is designed to help ensure that our dynamical principles do not rely on information that we cannot ascertain.

Several theorists have suggested the existence of this sort of locality desideratum. Einstein, for example, thought that the “principle of locality” was necessary for our ability to *investigate* the laws of nature empirically (Einstein, 1948). But of course, we now think that Einstein was spectacularly *wrong* about this; quantum mechanics appears to violate spatial locality, and it constitutes one of the most highly confirmed theories available, so Einstein can’t have been right that the principle of locality is a necessary presupposition in order for us to investigate the laws. The non-locality inherent in quantum mechanics was demonstrated by Bell’s theorem, along with the subsequent experimental violations of Bell inequalities.⁴⁶ The upshot of Bell’s theorem is this: according to quantum mechanics, what happens at one point in space can instantly affect what happens at another point, and in principle these two points can be extremely far apart. This appears to constitute a blatant violation of the spatial locality desideratum (even the revised one), and thus you might think that it spells trouble for my account of the Nomic Formula.

However, there are three points to note. First, and most basically, the extent to which we can design a theory that satisfies these various desiderata is dependent on how kind nature is. It may simply be the case that nature has been sufficiently unkind, such that we cannot rationally maintain this spatial locality desideratum. In fact, I think we really ought to *expect* not to be able to satisfy all of the desiderata simultaneously. Consider just the two desiderata of an informative dynamics and spatial locality. The latter desideratum directs us to use only local variables to

⁴⁶See Aspect et. al. (1981).

construct our dynamical principle(s). Thus it acts as a restriction on the class of variables we have available to us. On the other hand, the more variables we take into account, the more likely it is that we will be able to construct a dynamics that is deterministic and highly informative (albeit, also more complex). So there is already a fairly straightforward conflict between the desiderata of informative dynamics and spatial locality. Thus it should not be too surprising that we cannot satisfy both perfectly.

Second, even granting that nature precludes the possibility of a spatially local theory, it is still telling that our earlier theories (e.g. classical mechanics) *were* spatially local in the relevant respect. This suggests a picture whereby we initially tried to design theories that satisfied the locality desideratum, and it wasn't until our technology was sufficiently advanced that we discovered phenomena that caused problems for it. So the fact that quantum mechanics fails to satisfy the spatial locality desideratum is still compatible with the *existence* of such a desideratum. And the existence of such a desideratum is suggested, both by the a priori considerations I mentioned earlier (i.e. the fact that a radically non-local theory would be predictively useless to creatures like us), and the fact that our earlier theories *did* satisfy it.

Third, and finally, I should point out that even though quantum mechanics violates the spatial locality desideratum, it arguably does not do so *flagrantly*. In other words, it does not leave us in a situation of total predictive ignorance, because despite its non-local character, it is not *radically* non-local. We are still able to use quantum mechanics to make very reliable predictions even without knowing what is going on at very distant locations. This is because entanglement phenomena, which are the source of non-locality within the theory, usually play no

significant role in the behavior of macroscopic systems, and those are the kinds of systems that we often want to make predictions about.⁴⁷

4.4 Temporal Locality

A temporal locality desideratum arises for reasons very similar to the spatial locality desideratum just discussed. Let's return to our dynamical principle D describing the temporal evolution of system S . Again, D should ideally avoid requiring information that we cannot possibly ascertain. Now let's suppose that we've only just encountered system S , so we haven't been observing it for a while. Thus we don't know what it was doing two years ago, or a month ago, or even five minutes ago. If D now requires that kind of information about S 's history to make predictions about S 's future (i.e. if D is "non-Markovian"), we're going to be out of luck—we'll be unable to predict how S is going to behave. But if instead D only requires us to know the *current* state of S (i.e. if D is Markovian), then we're in better shape. In other words, it is useful for D to be *temporally* local: in order to make predictions about S 's future behavior, it only requires us to know the current state of S . The information required by Markovian principles is both *minimal* (we only need to know the system's state at a single instant) and the most *readily accessible* (we only need to know the system's current state). And it turns out that all of the dynamical laws postulated in physics (e.g. the Schrödinger equation and the various dynamical equations of classical mechanics) are temporally local in this respect.

⁴⁷The reason that entanglement does not play a significant role in the behaviors of macroscopic objects is controversial, and different interpretations of quantum mechanics explain it differently. On that note, the view I am developing here may afford an interesting explanation of the *very existence* of competing interpretations of quantum mechanics. Roughly, the thought is that quantum phenomena are so strange that they render it increasingly likely that conflicts will arise between these various desiderata. In other words, we are forced to choose which desiderata we want to satisfy, because we cannot satisfy all of them. Thus, different interpretations of quantum mechanics arise, satisfying subsets of these desiderata to varying degrees. Some, such as the Many Minds interpretation (Albert and Loewer, 1988), even preserve spatial locality. Obviously, this is an issue that deserves more exploration than I can give it here, but I think the basic idea is highly suggestive.

Note, again, that there is no reason that a Lewisian efficient summary would have to exhibit this kind of locality. From a God's-eye perspective, it may well be most efficient to summarize the universe by pointing out correlations between events at all sorts of temporal separations. Indeed, there is no more reason that an efficient summary would have to include principles that progress from earlier times to later times *than in the reverse*. On Lewis's view of laws, then, it would seem to be merely a stroke of luck—a convenient *accident*, as it were—that (a) there even *are* dynamical laws, and (b) these laws entail the dependence of later states on earlier states and not vice versa. But of course from *our* perspective, it is *no accident whatsoever* that the laws proceed from the past to the future. After all, the past (and present) is what we have information about; the future is what we want to make predictions about. Dynamical principles that proceed from the future to the past, or from the very distant past to the future, are predictively worthless to creatures like us. Thus the fact that all of the putative dynamical laws found in scientific practice are Markovian is, on this view, no accident.⁴⁸

4.5 Spatial, Temporal, and Rotational Symmetries

The spatial and temporal locality desiderata arise because we do not typically have information about spatially and temporally distant events. Another sort of information we often do not have is information that allows us to *locate* and *orient ourselves* within space and time. Or at the very least, we do not *begin* the project of constructing scientific theories with such information in hand (though we may eventually acquire it with the help of those very theories).

⁴⁸As I mentioned in §3, Woodward and Hall have argued that Lewis's BSA is likely to deem certain principles laws that intuitively are not. Similarly, Roberts has suggested that comparisons of theories on the basis of simplicity and strength cannot be found to appreciable degrees in scientific practice. Here we can begin to see *why* these objections get a foothold. Lewis has focused on the wrong sorts of desiderata in his Nomic Formula, thus engendering incorrect results about what counts as a law of nature.

Given this, it would be useful to have principles that do not require such information in order for us to be able to use them to make predictions. In other words, it would be useful to have principles that exhibit various symmetries, specifically spatial and temporal translational invariance, as well as rotational invariance. What these symmetries imply is that the principles can be applied the same way regardless of where in space you are located, when in time you are located, and how you are oriented with respect to the rest of the universe.

It is well known that the actual laws exhibit these various symmetries. Physicists regard this fact as highly significant for a variety of reasons, one of which is its connection to Noether's theorem. Noether's theorem states that such symmetries imply the existence of certain *conserved quantities*. In particular, spatial translational invariance implies the conservation of linear momentum, temporal translational invariance implies the conservation of energy, and rotational invariance implies the conservation of angular momentum. So one can view these various symmetries as implying the existence of the various conservation laws.⁴⁹

4.6 Predictively Useful Properties

In describing the preceding few desiderata, I have introduced them by saying that it would be useful if a given dynamical principle D did not require such-and-such unascertainable information in order to make predictions about system S . And remember that in the background here there is another desideratum, namely that D be as informative as possible—ideally *deterministic*. Now, I have suggested that the desiderata of spatial locality and informative dynamics can come into conflict: the latter places a restriction on our set of possible variables,

⁴⁹For an elementary yet illuminating presentation of Noether's theorem, see Feynman (1965, pp. 103-105). I should note that the existence of these symmetries does not by itself guarantee the existence of the corresponding conservation laws. Noether's theorem only applies to systems in which the principle of least action holds, so it's not as though the conservation laws simply *fall out* of the symmetry desiderata.

whereas the former is easier to achieve the larger the set of variables we have to work with. This issue also arises with temporal locality and the various symmetry desiderata: they place restrictions on the sorts of variables we can consider in the construction of our dynamics, thus making a deterministic, highly-informative dynamics harder to achieve.

Actually, however, we can avoid this conflict with a bit of philosophical subterfuge. For all I have said so far, a dynamical principle D can always be *forced* to simultaneously satisfy, for example, both the locality desiderata and the ideal of deterministic dynamics. To see this, let's look at a seemingly difficult example. Suppose that the system of interest S is a radioactive atom, and we want to predict whether it will decay in the next five minutes. Quantum mechanics teaches us that there is no way to tell whether or not it will decay: include any information you want about the history of the universe up to now, you still will not be able to tell whether the atom will decay within the next five minutes. Now, suppose the atom does decay, say at the four-minute mark. Once we detect the decay, we can posit a new property that the atom possessed four minutes ago, namely the property of being about to decay in four minutes. Nevermind that we couldn't figure this out ahead of time. We can still tell, after the fact, that it had this property four minutes ago. Thus if we *had* known everything about the current state of the atom four minutes ago, we would have been able to predict with certainty that it was going to decay.

In this way, we can construct a dynamics that respects the locality constraints and also achieves determinism. It is spatially local because whether the atom decays doesn't depend on anything far removed from the atom. It is temporally local because its decay at the four-minute mark depends on its having the property, four minutes ago, of being about to decay in four minutes. And it is deterministic: knowing everything about the atom four minutes ago, including

whether it has the property of being about to decay in four minutes, tells us with certainty whether it's going to decay.

The problem here should be obvious. We have no way of telling whether the atom has the property of being about to decay in four minutes except by *waiting for it to decay*. Thus by building this property into our dynamical principle describing the atom's temporal evolution, we have respected the locality desiderata as well as determinism, but we have still violated the guiding idea behind these desiderata, namely, that our dynamical principles should not crucially rely upon information that it is prohibitively difficult for us to ascertain. What we need, then, is some sort of constraint to prohibit this kind of maneuver.

What kind of constraint would do that? The problem here was that we posited a property that was impossible to ascertain until it became predictively useless: during the times when it would be useful to know, it is impossible to know, and once it can be known, it is no longer needed. So the constraint we need to add is one governing the *properties* that our principles are allowed to posit. What we want is for the properties to be ascertainable at a time when they can still be used to make predictions.

What sorts of properties does this rule out? It obviously rules out something like Lewis's property *F* (1983, pp. 41-42). The property *F* has, built into it, information that exhaustively details the entire history of the universe. It can be thought of as the property of "existing in a universe in which the following happens...[here follows a description of the entire history of the universe]," so it is possessed by everything in whatever universe it describes. Clearly we will never be in a position to discover that an entity has that property; it is predictively useless. And what's nice, in our case, is that we don't have to rule out *F* by fiat, as Lewis did. Lewis was worried about systematizations that could maximally satisfy the desiderata of simplicity and

strength simultaneously, by describing the universe as obeying the highly efficient and informative description, “ $(x)Fx$.” Thus he had to decree that systematizations must make reference only to *natural* properties, and *F* does not count as natural. However, Lewis notoriously took naturalness as a primitive, and so he could say nothing about what *makes* a given property natural, thus engendering all sorts of worries about whether we can know that we have succeeded in identifying the natural properties. In *our* case, on the other hand, no brute restriction to primitive naturalness is necessary. We merely require that the properties posited in our systematization be predictively useful to creatures like us. It’s obvious why such a restriction is necessary, and it rules out the sorts of properties that Lewis was worried about.

I should stress that this constraint on the allowable properties in our set of predictive principles does *not* rule out other sorts of “gerrymandered” or “gruesome” properties. It would not, for example, rule out the possibility that our principles should refer to the property of grueness rather than the property of greenness. For it can easily be ascertained *now* whether something is grue. So I am not claiming to have averted the New Riddle here. I am only claiming to have provided a very natural reason that *certain sorts* of gerrymandered properties—not *all* such properties—should not be expected to show up in our predictive principles. Of course, this opens up the possibility that we will still need to appeal to a primitive notion of naturalness, à la Lewis, in order to solve the New Riddle. That may well be the case, but it also may not. And given the possibility that we may not need to appeal to primitive naturalness elsewhere, I think it is fair to say that my theory’s ability to get by without making such an appeal still constitutes a distinct advantage over Lewis’s account.

It should go without saying, of course, that predicates such as Lewis’s ‘*F*,’ or such as ‘is about to decay in four minutes,’ do not show up in scientific practice. The existence of a

restriction to properties (or predicates) that are predictively useful is therefore confirmed by our actual practices.

4.7 Simplicity

The last desideratum of the Nomic Formula that I want to discuss is that of simplicity. Here I think it is helpful to approach the matter somewhat differently than I have above. In the foregoing, I have first been thinking in the abstract about what sorts of features creatures like us would look for in principles that we are going to use for making predictions, and then suggesting that these features are in fact possessed by the laws. In the case of simplicity, however, I think this strategy can be counterproductive. The practical benefits of simplicity are somewhat variegated, so trying to “reverse engineer” the preference for simplicity based on those benefits is a difficult task. So what I want to do here is proceed in the other direction: first, I will (briefly) argue that simplicity is used in the selection of physical theories about the laws of nature, and then examine some reasons that an appeal to simplicity in this manner might be helpful in making predictions.

I take it to be uncontestable that we do value simplicity in our scientific theories. To repeat a point from the previous chapter, consider how simplicity played a role in Gell-Mann’s introduction of the quark theory. In the 1964 paper in which he introduced the concept of the quark, Gell-Mann proposed two alternative explanations of the “eightfold way,” a highly successful classification system for hadrons. Both of these explanations involved the idea that hadrons were composed of more fundamental particles, but the first of them introduced a somewhat inelegant asymmetry in the posited properties of these more fundamental particles. In light of this, Gell-Mann suggested, “A simpler and more elegant scheme can be constructed if we

allow for non-integral values of the charges” (1964, p. 214). Thus the quark theory was born, along with the idea of fractional charge.

In *The Character of Physical Law*, Richard Feynman goes even further in his description of the role played by simplicity in the construction of physical theories:

One of the most important things [...] is to know when you are right. It is possible to know you are right [about a particular theory] way ahead of checking all the consequences [of that theory]. *You can recognize truth by its beauty and simplicity.* It is always easy when you have made a guess, and done two or three little calculations to make sure that it is not obviously wrong, to know that it is right. When you get it right, it is obvious that it is right—at least if you have any experience—because usually what happens is that more comes out than goes in. Your guess is, in fact, that something is very simple. If you cannot see immediately that it is wrong, and it is simpler than it was before, then it is right. The inexperienced, the crackpots, and people like that, make guesses that are simple, but you can immediately see that they are wrong, so that does not count. Others, the inexperienced students, make guesses that are very complicated, and it sort of looks as if it is all right, but I know it is not true because the truth always turns out to be simpler than you thought. (p. 171, emphasis mine)

Here Feynman is addressing the question of how we know when a given theory is correct. As he says, a theory obviously has to have the right sorts of observable consequences to be deemed correct. But just as obviously, this is not enough; it also should be simple. There may be very complex theories that get all of their observable consequences correct, but we know they are not true *because* of their complexity. To repeat: you can recognize the truth by its beauty and simplicity. Simplicity, then, is a guide to the laws of nature.

It is one thing to show that we value simplicity in the construction of physical theories. It is another thing to show why we value it. Now, what I am seeking here is not an epistemic justification of simplicity, but a pragmatic one. What I want to know is why simpler theories would be more valuable for pragmatic—specifically predictive—purposes. In other words, my central question is this: what is it about simpler hypotheses that makes them useful in making predictions?

There are fairly straightforward benefits of simplicity. For example, simpler hypotheses are easier to work with in a variety of ways. Not only do they tend to involve easier calculations, but there is heuristic value in adopting hypotheses that we are comfortable manipulating. Take, for example, the molecular theory of gases. This theory postulates that gases are made up of large numbers of particles, which essentially behave like miniature billiard balls in a constant state of motion. The familiarity and simplicity of the content of this theory allows us to make quick and rough predictions about how gases will behave without even doing any calculations. For example, consider someone unfamiliar with the macroscopic behaviors of gases, who only knows the simple postulates of the molecular theory of gases. Even to someone in such a position, it should not be surprising that heating a container of gas will increase the pressure inside the container. For it is not too difficult to reason that heating the gas will make the constituent molecules move faster, therefore causing more forceful and frequent collisions with the walls of the container. Of course, none of this is knowable a priori, but it does accord with our intuitive expectations. On the other hand, a theory with more abstruse theoretical postulates (that nevertheless makes the same predictions as the molecular theory of gases) is unlikely to lend itself to such easy predictions. In other words, it is harder to use a more complicated theory to get a rough idea of how the system of interest is going to behave.

Similarly, suppose a given theory makes some incorrect predictions, and we therefore decide that we are going to have to modify the theory. This task can be immensely easier if it is a simpler theory that is easy for us to grasp and manipulate. Since we are more easily able to understand how the constituent elements of a simpler theory combine to produce the various predictions, we are in a better position to know what elements might have to be changed when the predictions fail to come out as expected.

In short, simpler theories are easier to work with. They postulate theoretical structure that we know how to manipulate to generate rough-and-ready predictions without even doing any calculations, and they lend themselves to easier modifications in light of recalcitrant evidence.

Before ending this section, I want to recap the desiderata that I have claimed figure into the Nomic Formula. Contra Lewis, the Nomic Formula does not merely require us to balance simplicity and strength. Rather, it directs us to search for a set of principles with the following properties:

1. Informative dynamics
2. Applicability to a wide range of different systems
3. Spatial locality
4. Temporal locality
5. Spatial, temporal, and rotational symmetries
6. Properties and magnitudes that are predictively useful
7. Simplicity

I have argued that these properties are indeed found—to appreciable degrees—in the laws postulated in scientific practice. I have also argued that this should not be surprising. The rationale for searching for principles with these properties is that such principles are predictively useful to creatures like us. In other words, this is not just some random assortment of independently useful features. They all serve the common purpose of maximizing predictive utility. It is, then, entirely unmysterious why this collection of features—*precisely* this collection of features—should have coalesced under the concept of a law of nature.

5. Why are the Laws So Useful?

One of the benefits of the foregoing discussion is that we can use it to put pressure on alternative accounts of lawhood. I have already spent a good deal of time emphasizing the fact that Lewis's conception of the laws as efficient summaries from a God's-eye perspective is unlikely to produce laws with the sorts of pragmatic features discussed above. The fact that actual putative laws do have these features is therefore evidence against Lewis's view and in favor of my own. What I have not addressed, however, is how this fact can likewise be used to generate problems for nonreductionist theories.

Here is how to articulate the objection against those theories. First, let's highlight the following datum:

Nomic Predictive Utility (NPU): The actual putative laws of nature are remarkably predictively useful to creatures like us.

What I want to consider is how various nonreductive theories might attempt to explain NPU. I will argue that certain essential facts about nonreductive accounts of laws render them unable to provide a satisfying explanation of NPU. I will then explain why a reductive account like mine is better positioned to do so.

5.1 Predictive Utility and Nonreductionism

Let's look at a couple nonreductive theories of law to see how they might attempt to explain NPU. Take, for example, Armstrong's (1983) account. According to Armstrong, the laws are relations of necessitation between universals. If it is a law that all F s are G s, then according to his view a necessitation relation holds between the universals F and G . Armstrong writes this as $N(F, G)$. The pertinent question for us is this: Why, according to Armstrong's account, are the

laws predictively useful? Why, for example, should we expect necessitation relations between universals to tend to be spatially local? Why couldn't nomic necessitation relations hold between universals that get instantiated at great spatial distances?

Here, I think, Armstrong's theory is at a loss. Nothing in either his theory of laws nor his theory of universals seems to be able to account for the fact that the actual laws of nature exhibit spatial locality (to the extent that they do). On his view, one gets the sense that we must just be incredibly *lucky* that the laws have this useful feature.

Similar considerations apply to the desideratum of temporal locality. Now, it is somewhat tricky here, because Armstrong's laws do not have the form of dynamical equations. Neither $N(F, G)$ nor "All F s are G s" is straightforwardly translatable into a dynamical law.⁵⁰ Perhaps we could take F and G to describe temporal instants, and in that case a dynamical law could be forced into something resembling the form that Armstrong requires: "All F moments are succeeded by G moments." Of course, this would not be without its problems. Universals are not typically taken to describe temporal instants—they are thought to be such things as the fundamental *properties* posited by physics. Nevertheless, suppose we can find some way of extracting dynamical principles from Armstrongian laws. Why, in that case, should we expect those dynamics to be temporally local? No answers are forthcoming, on his view. There is no obvious reason why, say, future properties couldn't bear the nomic necessitation relation to past properties, or properties in the very distant past couldn't bear nomic necessitation relations to properties in the very distant future. Indeed, given all the theoretical alternatives on Armstrong's view, it should be *surprising* that the laws are temporally local. Again, if Armstrong's theory of laws is correct, it looks like we must just be incredibly *lucky*.

⁵⁰Maudlin (2007) takes this fact itself to be an objection to Armstrong's view.

It is important to head off a potential objection at this point. It might be tempting to reply on behalf of Armstrong in the following manner. “Look,” you might say, “I admit that, according to Armstrong’s account, we must be lucky. But *everyone* has to admit that. Nature has been, and continues to be, very kind to us. She could have been a lot more irregular, making it much more difficult to make accurate predictions. But on the whole she has been remarkably regular, and we therefore have to consider ourselves remarkably lucky *regardless* of what theory of laws we accept.”

This objection misses the point of my argument against Armstrong. I am not suggesting that Armstrong has particular trouble explaining why there have been *regularities in the particular matters of fact* that are amenable to our predictive endeavors. What I am suggesting, rather, is that Armstrong has trouble explaining why *the laws* are amenable to our predictive endeavors. For on his view, there are a *wide variety* of ways the laws could be that are compatible with the regularities in the particular matters of fact being exactly as they are. And on many of these ways the laws could be (according to him), they would not be predictively useful. Therefore, what his view has particular difficulty explaining is why the *laws*, not the regularities, are amenable to our predictive interests.

This problem is not unique to Armstrong’s theory. For example, Maudlin (2007) takes the laws to be fundamental elements of the universe that govern its temporal evolution: they generate successive states of the universe from its current state. Now such a view has, built into it, the temporal locality that I have emphasized is so crucial for the laws’ predictive utility. So here it might seem that Maudlin is in a better position than, say, Armstrong. But of course, his conception of laws does not have other features built into it, such as the laws’ spatial locality, or their various symmetries, or their applicability to a wide range of systems (as opposed to just the

total state of the universe at any instant). Like Armstrong, then, he must regard the fact that the laws possess these features as a happy accident.

Furthermore, it is not all that clear that Maudlin has improved his position by building temporal locality into the metaphysical nature of laws. Since he takes laws as ontologically primitive, I doubt that he can offer a satisfying explanation of why they exhibit this characteristic. Of course, he might just argue that the laws exhibit temporal locality *by metaphysical necessity*, since it's just *part of what it is to be a law*, on his view. But then his burden is to explain why metaphysical necessity happens to be so amenable to the predictive desires of creatures like us. Faced with such a burden, Maudlin could suggest that metaphysical necessity simply does not admit of any further explanation. And that may well be the case. But it does not change the fact that the laws are temporally local, and this fact seems to deserve a better explanation than just an appeal to brute metaphysical necessity. By making such an appeal, the ability of Maudlin's theory to furnish a satisfying explanation of the laws' temporal locality—and indeed their predictive utility more generally—would be greatly diminished.

It is not just Armstrong's and Maudlin's theories that are open to this sort of objection. Indeed, any theory that denies the supervenience of the laws on the particular matters of fact will be vulnerable to these types of worries. Lange (2000, p. 51) gives a helpful characterization of the failure of supervenience: "I tend to picture the [facts of the form 'it is a law that *s*' and 'it is not a law that *s*'] as having been sprinkled like powdered sugar over the doughy surface of the non-nomic facts." There are, of course, many ways to sprinkle powdered sugar over a ball of dough; the very same ball of dough could be the recipient of *innumerable* different sprinklings. Now, on any view that treats the non-nomic facts like a ball of dough and the laws like powdered sugar, the question is going to arise why the powdered sugar has fallen in such a convenient

pattern for us denizens of the dough. Granted, one could simply reply that there is no explanation; the powdered sugar just fell that way, end of story. But is that really plausible? If I sprinkle some powdered sugar on a ball of dough and it falls in the pattern of Mike Tomlin hoisting the Lombardi trophy, it would be *very* difficult for me to believe that this is just happenstance. I would first have to search extensively for an explanation; only if I could not find one would I ever accept that it was coincidence, and probably not even then.⁵¹

5.2 Predictive Utility and Reductionism

I now want to explain why the predictive utility of the laws is much easier to accommodate on a reductionist metaphysics. According to the reductionist, laws are certain sorts of patterns in the phenomena. *What sorts* of patterns? Well, different reductionist theories offer different criteria for picking out the laws among all the patterns that obtain. But it is perfectly open to the reductionist to claim that the laws of nature are those patterns in the phenomena that are most predictively useful to creatures like us. By saying this, the reductionist is able to provide an explanation of why the laws have various features that make them predictively useful. Their predictive utility, according to the reductionist, is no accident; the laws are selected and distinguished *based on* their predictive utility. What puts the reductionist in a better position than the nonreductionist here is that, according to the reductionist, there is an *abundance* of patterns that are, metaphysically speaking, just like the laws. The presence of certain features in the laws can thus be explained as the result of a selection process operating on this plurality of patterns. The nonreductionist has no analogous move available. According to her, the laws are

⁵¹Incidentally, I would not be at all satisfied if someone told me that the powdered sugar fell in just this pattern because of *metaphysical necessity*. Would that metaphysical necessity were a Steelers fan.

metaphysically *sui generis*, so she cannot explain why they have certain features by appeal to a selection process operating on a plurality of similar sorts of entities.

Now, you might think that there is a potential reply open to the nonreductionist here. Perhaps the nonreductionist can attempt a similar maneuver, and couple it with an anthropic argument. To see what I mean, let's focus on a particular nonreductionist theory. Here I will examine Armstrong's theory, merely for ease of exposition.⁵²

We have seen that the reductionist has available to her a plurality of entities that are metaphysically just like the laws (on her view), among which the laws could be distinguished by pragmatic constraints designed to pick out predictively useful principles. Suppose Armstrong tries to avail himself of a similar maneuver. According to his theory, laws are necessitation relations between universals. Of course, there may be *all sorts* of necessitation relations between universals; these relations may be local, nonlocal, symmetric, asymmetric, etc. But which of those necessitation relations are creatures like us most likely to be interested in? Quite plainly, we are most likely to be interested in the ones that are predictively useful to creatures like us. So of course the laws *that we know about* are going to be predictively useful to creatures like us. This is not because predictive utility is part of what it is to be a law, but merely because *we are most likely to look for* those necessitation relations between universals that are predictively useful. The existence of such relations does not preclude the possibility that there are other necessitation relations between universals—other laws—that are not at all predictively useful to creatures like us; it's just that we haven't found such relations, because we haven't looked for them.

⁵²A similar argument, and a similar rebuttal, could be given for other nonreductionist theories.

The problem with this response is that we tend to think, *pace* Cartwright (1999), that the actual putative laws are *complete*, in that they “subsume” all the events that occur in the universe. For example, the desideratum of wide applicability for the laws (§4.2) pulls in this direction. And insofar as the laws are *not* complete, then there are domains where they cannot be used to make accurate predictions. Should we discover such domains, we would surely seek to modify our current laws so that they *do* apply there.

If the actual putative laws are complete (or nearly so), then the above response on behalf of Armstrong is going to imply a very peculiar and widespread *overdetermination* of the particular matters of fact by the laws. To see this, consider some particular event *e* that we want to explain by appeal to the laws. If the actual putative laws that we have discovered subsume every event in the universe, then we should be able to explain why *e* occurs by appealing to such laws. But we should *also* be able to explain why *e* occurs by appealing to other, disjoint sets of laws. For the above response requires that there are a lot more laws out there, many of which are not predictively useful. So while we may not be able to use these other laws to predict that *e* is going to occur, we *should* be able to use them to explain why it does occur.

This strikes me as a very odd situation. For one, these extra laws seem theoretically otiose; there is no need to posit them other than to give the above explanation of why the actual putative laws that we know about are so predictively useful. They are not required to account for any events in the history of the universe, for the actual putative laws that we know about already do so. Furthermore, what ensures that these other laws agree with the known laws about all the particular matters of fact? That is, for every event *e* in the history of the universe, these disjoint sets of laws (of which there are who knows how many) agree with the known laws about how those events play out. And this has the flavor of a massive conspiracy. How have all these

disjoint sets of laws made sure that they will always have the same implications for every event that ever occurs? Here we seem pushed to posit some sort of *meta-law* that constrains all of these first-order laws to make sure they agree on exactly how they are going to overdetermine all of the particular events in the universe's history.

The absurdity of this response should now be manifest. In short, its metaphysical extravagance seems unjustified, especially given that a much more satisfying explanation of the laws' predictive utility is readily available.

5.3 Why the LOPP is Misleading

Before concluding, I want to address one more objection. There is a natural worry about an account of laws that, like mine, combines (1) some sort of regularity conception of lawhood with (2) an emphasis on predictive utility.⁵³ The thought, roughly, is this. According to my account, the laws of nature are the outputs of applying the Nomic Formula (i.e. a formula designed to generate the most predictively useful principles) to the Humean base (i.e. the totality of facts about macroscopic phenomena). But note that if we *had* all the information about the macroscopic phenomena, there would be *no point* in looking for predictively useful principles. For there would be nothing left to make predictions *about*. And so the thought is that the very idea of a law of nature, on my view, is internally confused.

We can sharpen this objection by appeal to the LOPP from §2. According to my theory, we are supposed to give the LOPP the totality of facts about macroscopic phenomena. She then applies the Nomic Formula to these facts, and out come the laws and fundamentalia. Now as I argued in §4, the Nomic Formula is designed to output principles that are predictively useful. So

⁵³Something in the neighborhood of this worry is articulated by Woodward (2014, pp. 115-116).

the problem is this: there is *no reason* for the LOPP to apply the Nomic Formula to the totality of macroscopic phenomena, for there is nothing left for her to make predictions about. From her perspective, the architecture of my theory of laws simply does not make sense.

This objection, I think, is exactly right. The only thing wrong with it is that it is not an objection. What it shows us, rather, is that from a God's-eye perspective, it does not make sense to be concerned about the laws of nature. The laws are not useful to a god, or even to a LOPP. Any being that knows how all of the macroscopic phenomena play out has no reason to search for the laws of nature. But of course, creatures like us *don't* know how all the macroscopic phenomena play out. For us there *is* reason to try to discover the laws of nature. So from our perspective, there is nothing incoherent, nothing irrational, about applying the Nomic Formula to what we know about the macroscopic phenomena.⁵⁴

In hindsight, this actually should not be too surprising of a result. Remember, on any reductionist view, there is a *prima facie* challenge: we must be able to explain what makes the laws of nature “distinctively appropriate targets for scientific inquiry,” as Hall puts it. In other words, we must be able to explain why we should be interested in discovering the laws. The explanation suggested by Hall, and developed here, is that discovering the laws has pragmatic benefits for creatures like us. They are convenient systematizations that we have pragmatic reasons to care about. This is all perfectly compatible with the laws *not* being interesting to creatures who are in very different positions than we are. The *data*, as it were, show that *we* are

⁵⁴C.f. Roberts (2008, pp. 343-346). Roberts's account of laws also has the consequence that a being who knows all the details about particular matters of fact would not care about the laws. While the technical details giving rise to this consequence are very different for his theory, he regards it as unobjectionable for similar reasons: “God's Own Theory has no laws, for laws are a feature of theories used by subjects who know only a tiny fraction of the contents of space-time, and need to rely on measurements whenever they want to find out more” (p. 344).

interested in the laws. Nothing tells us that a god, or even an oracle, would likewise have to be interested in them.

So thinking of this view in terms of the LOPP can actually be misleading. There are crucial respects in which her position is different from our own, and these differences affect the very motivations for trying to discover the laws in the first place. She is more like the person navigating a maze on a piece of paper, rather than the person inside a literal maze. Indeed, John Roberts has suggested that a better way to think of her would be, not as an ideal physicist, but as something like a guardian angel. She can see everything, and she summarizes it in a way that is maximally predictively useful for us. We could even explicate the theory by reimagining Lange's imaginary conversation with God:

You: Lord, I have no idea what I'm doing here. Could you help me out?

God: Sure. Right now, there's a particle in state Ψ_1 , and on the other side of the universe there's a particle in state Ψ_2 , and when the second particle moves in a trajectory described by this fancy equation, the first particle will move in a trajectory described by this other fancy equation...

You (frustrated): Lord, I can't see the other side of the universe.

God: Good point. In that case, I'll give you a set of principles that would be maximally predictively useful to someone in your position.

You: Do tell...

6. Conclusion: Why We Care about the Laws

In this chapter I have developed an account of the Nomic Formula, the second main component of my view of laws. The Nomic Formula is supposed to be the collection of

desiderata that scientists themselves try to satisfy when they investigate the laws, and it is designed to output principles that are maximally predictively useful to creatures like us. Conceiving of the Nomic Formula in this way has several advantages. First, it ends up producing principles that exhibit many of the formal features possessed by actual putative laws of nature, both past and present. I take this to be evidence for its correctness. Second, it makes sense of scientists' epistemic practices. Essentially, it elevates the epistemic standards used by scientists to standards that are *constitutive* of lawhood. What better way to vindicate those very standards? Third, and most significantly, it makes sense of *why scientists are interested in the laws in the first place*. The point of trying to discover the laws of nature is that, by doing so, we arrive at principles that are of great predictive utility to creatures like us.

Thus, by structuring the view this way, we vindicate scientists' epistemic standards on epistemic grounds, *and also on pragmatic ones*. In other words, this view allows us to explain, not only why our epistemic standards are a reliable guide to the laws, but also why we should be interested in implementing those standards in the first place. Furthermore, insofar as we take discovering the laws of nature to be one of the central tasks of science, we also gain an illuminating explanation of the reasons that creatures like us would be likely to invest so heavily in this sort of enterprise.

This completes the core elements of my view of laws. The task is now to argue that, by conceiving of the laws in this way, we are able to make sense of a number of other important features that we traditionally take them to possess. Specifically, in the next chapter, I turn my attention to the close relationship between laws and counterfactuals. My aim is to account for this relationship using the theory of laws that I have just developed.

CHAPTER 4: WHY DO THE LAWS SUPPORT COUNTERFACTUALS?

1. Introduction

When we engage in counterfactual reasoning, we tend to suppose that the actual laws of nature still hold. For example: if I had gotten up later this morning, no laws of nature would have been violated. And if the moon did not exist, there would be no neap tides, because the law of gravity would still hold. There is little doubt about the fact that we reason like this, and many accounts of laws and counterfactuals have been designed to try to accommodate it.⁵⁵ But *why* do we reason like this?

This chapter aims to answer that question in a way that is consonant with the theory of laws I have developed in Chapters 1 and 2. I proceed as follows. In §2, I review and endorse Maudlin's (2007) "Altered States Recipe" for the evaluation of counterfactuals. It turns out that if the Altered States Recipe is correct, the laws will be held fixed in the evaluation of counterfactuals. The question then is why that Recipe should be correct. In §3 I will offer an answer to that question. Roughly, my answer is that creatures like us use counterfactual reasoning to try to ascertain facts about the *actual* world, and in doing so, we are greatly aided by the predictive power of the laws of nature. So the laws are held fixed in counterfactual reasoning because doing so helps us achieve our *goal* in counterfactual reasoning: discovering facts about actuality.

⁵⁵See, for example, Lewis (1973, 1979), Lange (2000, 2009), Maudlin (2007), Goodman (1983), and Roberts (2008).

This conclusion is compatible with a variety of different theories of laws, but it is particularly helpful to the Best Predictive System Account. In §4, I explain why this is the case. In particular, I argue that my explanation of counterfactual reasoning does two things for the Humean: (1) it accounts for the laws' counterfactual resilience in a way that is compatible with a Humean metaphysics, and (2) it helps to explain why we may have anti-Humean *intuitions* even if some form of Humeanism is correct. Finally, I conclude in §5.

One caveat before we proceed. So far I have spoken of the laws' support of counterfactuals, and the counterfactual resilience of the laws, and I have said that the laws are held fixed in the evaluation of counterfactuals. In saying all of this, I mean to leave open whether it is the *lawhood* of the laws, or merely their *truth*, that is held fixed, resilient, etc. Many theorists have maintained the former, i.e. that the laws are *still laws* under any counterfactual supposition that is consistent with all of them. This claim is often referred to as "Nomic Preservation" in the literature.⁵⁶ I do not mean to take on this further claim, though I do think that my proffered argument is consistent with it.

2. The Altered States Recipe

Maudlin (2007, pp. 21-34) provides an elegant recipe for the evaluation of counterfactuals.⁵⁷ He calls this a "recipe," and not a "theory," because he does not want to commit himself to the claim that all counterfactuals must always be evaluated in this manner. Rather, he maintains (and I agree with him) that this recipe applies to at least a "large number of types of counterfactuals," and that we tend to be most confident in the correctness of our counterfactual assertions when

⁵⁶Supporters of Nomic Preservation include Lange (2000, 2009), Goodman (1983), Chisholm (1955), and Jackson (1977). Roberts (2008) maintains that Nomic Preservation is true in all scientific contexts, though possibly not in some other contexts.

⁵⁷See also Paul and Hall (2013, pp. 47-53).

they agree with the results of this recipe. While reviewing the recipe, I am going to highlight some salient features, to which we will return in §3.

Here is the recipe. Consider a counterfactual of the form, “If A had been the case, then B would have been the case,” symbolized $A \Rightarrow B$. Step 1 is just to choose a time slice of actual history.⁵⁸ All physical magnitudes will take some value on this time slice. Which time slice we must choose is a matter of context, and sometimes it will not matter exactly which one we pick, as long as all of the time slices in a reasonable window give us the same results. If it *does* matter which one we pick, and context does not narrow down our choice enough, then the truth-value of the counterfactual will be indeterminate.

Step 2 is then to modify just enough physical magnitudes in our selected time slice to make A true. (If A was true in the actual world, then we modify nothing.) Crucially, we do not modify any more than we have to: we want the altered state to be as much like the actual state as possible. Let us therefore say that the Altered States Recipe is *conservative*.⁵⁹

Sometimes there will be multiple ways to make A true—i.e., there will be a variety of different possible modifications of physical magnitudes that would all achieve a situation in which A obtains—and in such cases we rely on context to determine which physical magnitudes we have to modify. And again, if context is insufficient, then the truth-value of the counterfactual may be indeterminate.

Step 3, then, is to plug our modified time slice into the laws of nature and see what model(s) they generate as a result. If the only models generated are ones in which B occurs, then

⁵⁸Maudlin talks of selecting a Cauchy surface rather than a time slice, but nothing crucial will hang on the difference here.

⁵⁹David Lewis’s (1973, 1979) possible worlds semantics for counterfactuals is also conservative, in the sense that we are directed to consider an A -world that is as similar as possible to the actual world (at least before the time of the antecedent).

the counterfactual is true. If not, the counterfactual is false. Let us therefore say that the Altered States Recipe is *nomically guided*.

Consider an example: “If the moon did not exist, there would be no neap tides.” To evaluate this counterfactual according to the Altered States Recipe, we first have to choose a time slice. This particular counterfactual seems to give us very wide latitude about which time to choose, and it will turn out not to matter much. Suppose we pick a time slice from the year 1800. Now we modify the physical magnitudes on that time slice to remove the moon. This will involve deleting all the particles in the moon’s location and replacing them with empty space (presumably). It may also involve other changes, but these need not concern us here. Of course, we don’t make changes if we don’t have to. We don’t change the mass of the sun, for example, nor do we add a moon-sized Death Star in place of the original moon. We then plug this altered state into the laws and let it evolve. If there are still neap tides, the counterfactual is false. If there are no neap tides, the counterfactual is true. (And of course, it *is* true.)

This is the Altered States Recipe in its essentials. However, there is a complication for cases involving indeterministic laws, and it will be worth considering how these sorts of cases are handled. Suppose that coin flips are a fundamentally indeterministic process. Also suppose that I am a billionaire, and I am about to flip a coin (fair or not). I offer you the following bet: if you correctly call what side the coin will land on, I will pay you \$1 billion, but if you are wrong, you pay me \$1 billion. Naturally, you decline the bet. I flip the coin. It lands heads. You think to yourself, “I made the smart choice, but shoot, if only I had bet on heads, I would be rich now.”

You’re right, of course: if you had bet on heads, you would now be \$1 billion richer. But as it stands, the Altered States Recipe does not deliver this result. Suppose we select a state before I flip the coin, and before you decline the bet. We change just enough of the physical

magnitudes in that state to make it true that you bet on heads. We then run that state forward, according to the indeterministic laws. These laws generate two sets of models, one on which the coin lands heads, and one on which the coin lands tails. In the “Heads” models, you are \$1 billion richer, but in the “Tails” models, you are \$1 billion poorer. Since there is no unambiguous result, the counterfactual “If I had bet on heads, I would be \$1 billion richer” is false.

There are some philosophers who agree with this result. For example, Redhead (1987) argues that if you had bet on heads, you might have been right, but you might also have been wrong, even though the coin *actually* landed heads. Lewis’s (1979) account of counterfactuals also has this consequence (though this is merely because he did not design it to accommodate indeterministic cases). I think this result is clearly wrong. Given that the coin actually landed heads, if you had bet on heads, you would have won. Maudlin agrees, noting that Redhead’s account conflicts with David Mermin’s (1990) Strong Baseball Principle, namely that I cannot affect the outcome of a baseball game being played in Pittsburgh by fiddling with my television here in Massachusetts. (The Pirates lose *regardless* of whether I turn the television off or not.) This is supposed to be true even if the fundamental laws of nature are irreducibly probabilistic. But Redhead would have to say that, if the laws are fundamentally probabilistic, then what I do with my television might indeed affect the outcome of the game, for if we alter my interactions with the television in the minutes before the game, and then plug that altered state into the indeterministic laws, they will generate a set of models, some of which differ in the outcome of the game. This certainly stokes my superstitious side (“If I hadn’t watched, they might have won...”), but I think it is untenable, and will assume as much here.⁶⁰

⁶⁰Actually, Marc Lange points out that Redhead need not be committed to the claim that my actions *affect* the outcome of the game. That would depend on the further thesis that causation is to be analyzed in terms of counterfactual dependence. So strictly speaking, all that Redhead’s account commits him to is the claim that if I had

Back to the coin flip. If indeed you would have won, had you bet on heads, then we need to modify the Altered States Recipe to deliver this result. The natural way to do so is to rule out some of the models generated by indeterministic laws. But notice that we cannot just rule out *all* models that disagree with the actual outcome of heads. If, in the process of flipping the coin, I had raised my hand up higher before letting it go, then it is *not* guaranteed that the coin would still have come up heads. So *some* counterfactuals hold fixed the outcome of indeterministic processes, but some do not. Specifically, as Maudlin puts it, “if there had been a different causal process leading to the result, we might have gotten a different result. If the process is unchanged in the counterfactual, so should the result be” (p. 29). In the case of the coin, if the causal process that leads to the flipping of the coin is unchanged by the counterfactual supposition, then the outcome of the coin flip should also be unchanged. But if the causal process *is* changed, then the outcome might be different.

Maudlin sketches a solution to this problem that requires the concepts of “infected” and “uninfected” magnitudes. Recall that in Step 2 of the Altered States Recipe, we are directed to alter some of the physical magnitudes in our chosen time slice in order to get the antecedent of the counterfactual to come out true. Any of the magnitudes that we modify here are deemed “infected.” As Maudlin says, any stochastic process can be viewed as a set of evolving physical magnitudes, and the laws tell us both *how* later magnitudes depend on earlier ones, and *which* later magnitudes depend on which earlier ones. Or at least, the laws tell us this if they have the features recommended by the Best Predictive System Account. In particular, they need to have extensive dynamical implications for the system(s) in question, and they need to treat those systems as at least quasi-enclosed—else the distinction between infected and uninfected

turned the television off, the Pirates might have won. This may be slightly less unreasonable than the claim that my actions affect the outcome of the game. But only slightly.

magnitudes would quickly break down. Any magnitude that depends on an earlier infected magnitude is *also* infected, but no other magnitudes are infected.

The prescription of the Altered States Recipe is then as follows. Steps 1 and 2 remain the same. In Step 3, we again let the altered state evolve according to the laws. If the laws are indeterministic, they will generate a variety of different models based on that initial state. We are then directed to compare these models to the actual history: any models that diverge from the actual history based solely on *uninfected* magnitudes are disregarded. On the other hand, models that diverge, and whose divergence is traceable to *infected* magnitudes, cannot be disregarded. The counterfactual is true iff the consequent obtains in all of the models that are not disregarded.

The net effect of this is to ensure that chance events that are *not* the result of our modifications (to the selected state in Step 2) are held fixed to their actual outcomes. So what would have happened if you had bet on heads, according to this recipe? Here we select our initial state, and modify it just enough to make you take the bet on heads, but no more. Specifically, we do *not* change anything about how *I* am positioned, about to flip the coin. We then let this state evolve according to the laws. Two sets of models are generated, one on which the coin lands heads, and one on which it lands tails. But, if everything goes as expected, the infected magnitudes constituting your taking the bet on heads do not spread to the flipping of the coin, which is therefore comprised solely of uninfected magnitudes. Consequently, we disregard the Tails models, because they are the result of uninfected magnitudes and they disagree with the actual history. Thus we get the desired result: all of the relevant models are ones in which the coin lands heads, and you have won \$1 billion.

The prescription here can be put rather more straightforwardly, without regard to infected and uninfected magnitudes: in evaluating a counterfactual $A \Rightarrow B$, any magnitudes that are

causally independent of A must be fixed to the same value that they take in the actual world, even if they are the result of indeterministic processes. Roughly, the thought is that if, in hindsight, we know that some chance event C occurred, and we know that C 's occurrence was causally independent of the occurrence or non-occurrence of A , then we should hold fixed C 's occurrence in evaluating $A \Rightarrow B$. Let us therefore say that the Altered States Recipe uses *hindsight*.⁶¹

Insofar as the Altered States Recipe is correct, the laws are held fixed in the evaluation of counterfactuals. We can see this by checking each step of the Recipe. At Step 1, we select a state of actual history, and of course the actual history is compatible with the actual laws of nature. At Step 2, then, we alter this state to make the antecedent of the counterfactual true. In order for us to be able to plug this state into the laws at Step 3, the altered state must itself be compatible with the laws.⁶² Then at Step 3, we evolve that altered state forward (or backward, as the case may be) according to the laws, while ignoring any models that disagree with the actual history about the occurrence of chance events that are causally independent of the truth of the antecedent. Since we are only directed to *ignore* certain models generated by the laws, and never to add any models in, there is no possibility that we ever consider a history that violates the laws.

So at no step of the Recipe is there any possibility of introducing a violation of the actual laws. In short, as Maudlin puts it, “We construct the counterfactual situation by means of the laws, so the laws must hold” (p. 34).⁶³

⁶¹I borrow this terminology, and this way of articulating the method, from Edgington (2004).

⁶²Of course, some counterfactual suppositions are themselves inconsistent with the laws, so there is no way to create an altered state that both (1) makes the antecedent true, and (2) is consistent with the actual laws. It is an important question how to handle such “counterlegals,” but that is beyond the scope of this chapter.

⁶³Note, however, that it is left open here whether the laws are still *laws* in the counterfactual situation, or whether they are merely still *true*. The laws figure into the recipe by allowing us to construct models of counterfactual situations that follow, via the laws, from the altered state. All we *need* to assume for this to work is that the

I agree with Maudlin that we often evaluate counterfactuals in accord with the Altered States Recipe. I also agree that we tend to be more confident in the correctness of our counterfactual assertions when they are evaluated by this method. So I am going to assume that the Altered States Recipe is correct, at least for a very broad range of counterfactuals. The question now is *why* it is correct. Why should we evaluate counterfactuals in accord with this method? In particular, why should our counterfactual reasoning be *conservative* and *nomically guided*, and why should it use *hindsight*?

These questions get different answers on different theories of laws and counterfactuals. Maudlin, for example, seems to regard it simply as a brute fact that the nomic structure determines the counterfactual structure in accord with the Altered States Recipe. In other words, the Recipe is right because it correctly describes some genuine, *sui generis* counterfactual structure that exists independently of us and our interests. Viewing the counterfactual structure as a *sui generis* entity gives Maudlin straightforward answers to the questions of why our counterfactual reasoning is conservative and nomically guided, and why it uses hindsight: if it didn't have these features, it would deliver incorrect descriptions of the objective, independent counterfactual structure. Of course, these answers are straightforward, on Maudlin's theory, at the cost of also being *trivial*. Maudlin can provide no deeper explanation of why our counterfactual reasoning exhibits these features aside from the fact that it *must*, in order to generate correct descriptions of counterfactual reality.

Other theories of laws generate less trivial answers to some of these questions. Lange's (2009) theory, for example, regards the laws' resilience under counterfactuals as a defining feature of the laws. Lange accomplishes this by first explicating the general notion of sub-nomic

regularities described by the laws still obtain, but that is consistent with them not being laws. So the correctness of the Altered States Recipe does not settle the issue of Nomic Preservation.

counterfactual stability. A set of propositions is *counterfactually stable* if all of its members would still have been true under any counterfactual supposition consistent with that set. A set of propositions is sub-nomic if they express only sub-nomic facts, i.e. facts that do not imply anything of the form “it is a law that...” or “it is not a law that...” Lange thus identifies the laws as the members of the largest non-maximal sub-nomically stable set. The laws, on Lange’s view, are what they are because they are determined by the counterfactuals. Thus the fact that our counterfactual reasoning is nomically guided is a direct consequence of *what it is to be a law*.

However, Lange also takes subjunctive facts as primitive. So whereas he is able to provide an illuminating answer to the question of why our counterfactual reasoning is nomically guided, his answers to why it is conservative and uses hindsight would appear to be the same as Maudlin’s: if our counterfactual reasoning didn’t have these features, then it would generate incorrect descriptions of the subjunctive facts. End of story.⁶⁴

By contrast, in the next section I develop an alternative explanation of why our counterfactual reasoning has these three features. In particular, I argue that creatures like us would find it *useful* to reason about counterfactuals in accord with the Altered States Recipe, and thus I provide nontrivial reasons that our counterfactual reasoning is conservative, nomically guided, and uses hindsight. To make this argument, I am going to have to explain why creatures like us are interested in counterfactuals in the first place.

⁶⁴At least, it would seem natural for Lange to provide such answers, though I should note that he does not explicitly address the questions of why counterfactual reasoning is conservative and uses hindsight.

3. Counterfactuals as a Guide to Actuality

In this section, I will argue that we engage in counterfactual reasoning with the goal of figuring out facts about actuality.⁶⁵ This goal makes sense of why our counterfactual reasoning is conservative and nomically guided, and why it uses hindsight.

To begin, consider a familiar equation:

$$P(H|E) = \frac{P(E|H)P(H)}{P(E|H)P(H) + P(E|\neg H)P(\neg H)}$$

Bayes theorem gives expression to a number of truisms about how a given piece of evidence E bears on a hypothesis H . I want to focus on one of them. For $P(H|E)$ to be high, we want $P(E|H)$ to be higher than $P(E|\neg H)$ to be low.⁶⁶ In other words, a given piece of evidence speaks strongly in favor of a hypothesis when the only plausible way for that evidence to come about is for that hypothesis to be true. If the butler would probably have used the candlestick, but everyone else would probably have used either the gun or the knife, then the fact that Mr. Boddy was killed with the candlestick strongly supports the hypothesis that the butler did it.

Sometimes we are able to anticipate the possibility of some evidence in advance. In such cases, we reason roughly as follows: “The probability that E will occur, if H is true, is very high. And the probability that E will occur, if H is *not* true, is very low.” So when we find that E occurs, we have a straightforward inference to the truth of H . Indeed, if we were logically omniscient, we could anticipate every possible piece of evidence, and how likely each would be on H and $\neg H$. So whenever we encountered a new piece of evidence, all we would have to do is straightforward conditionalization to update our credence in H .

⁶⁵Cf. Edgington (2004) for a similar proposal.

⁶⁶Of course, we also want $P(H)$ to be high, and $P(\neg H)$ to be low, but consideration of these priors is not germane to our present purposes.

But of course, creatures like us are not logically omniscient. Much of the evidence we encounter, we did not anticipate the possibility of encountering ahead of time: we never evaluated, nor even thought of the *possibility* of evaluating, $P(E)$, $P(E|H)$, or $P(E|\neg H)$. (For example, I bet that you did not anticipate the possibility of reading exactly *this* sentence right now, but here you are, reading it anyway.) When we encounter evidence that we did not anticipate, it no longer makes sense to ask ourselves, “Will this evidence occur if H is true?”, for the evidence has already occurred. Instead, to evaluate its bearing on a given hypothesis H , we have to ask, “*Would* this evidence have occurred (or *how likely* would it have been to occur) if H were true? What about if H were false? In other words, we have to evaluate subjunctive conditionals, at least one of which we know must express a counterfactual.

So the suggestion is that we use counterfactual reasoning to figure out how our evidence bears on the truth of various hypotheses of interest when we did not anticipate the possibility of getting that evidence in advance. In a word, counterfactual reasoning is our solution to the problem of unanticipated evidence.⁶⁷ If this is correct, it explains why counterfactual reasoning is conservative and nomically guided, and why it uses hindsight. I’ll consider these in turn.

First, why does this use of counterfactual reasoning require it to be conservative? Suppose we are investigating the scene of a car accident, and we are trying to figure out how fast the car was going. The driver claims they were going 30 miles per hour before the accident, but

⁶⁷Note that this is distinct from the so-called “problem of old evidence” (as originally posed in Glymour (1980)). The problem here is not that we have a newly-formulated hypothesis designed to account for previously-known evidence. It is rather that we have newly-acquired, unanticipated evidence, and we want to figure out which hypothesis rendered that evidence most likely.

Nevertheless, the problems are similar in that they both require us to evaluate the relationship between hypotheses and evidence in contexts where we do not want to take for granted that the evidence has actually occurred. The difference, as I see it, is that the focus here is not on what one’s subjective probability of the evidence would have been if one did not already observe it (which can often be hard to answer for a variety of reasons), but rather on an objective relationship between hypothesis and evidence—precisely that relationship that is captured by the Altered States Recipe. Or so I am arguing.

seasoned investigators like ourselves do not accept testimony as a form of ultimate evidence. Instead, we decide to examine the length of the tire tracks leading to the crash. We do so, and ask ourselves: “Would the tire tracks have been this long if the car had been traveling 30 miles per hour?” To evaluate this, we first select a time slice before the accident, and construct it so that the car is traveling 30 miles per hour toward the scene of the accident. We keep everything as similar to actuality as possible: we don’t change the weather, nor the driver of the car, nor the age of the tires that were on the car, nor the temperature of the road, nor the weight of the car, etc. Why not? Because our goal is to use the tire tracks to figure out what speed the car was *actually* traveling, and to do that we need everything else that could have affected the tire tracks to remain the same. If we altered these aspects of the time slice, if we departed from actuality in these respects, then we wouldn’t get a good read on the actual speed of the car.

Why is our counterfactual reasoning nomically guided? Once we’ve selected the time slice, and fixed the speed of the car, we use the laws to evolve that state forward, paying specific attention to how long the tire tracks end up being. Why do we use the laws to do this? Because, if the Best Predictive System Account is correct, the laws are remarkably effective ampliative inference machines: they express connections between distinct existences, allowing us to infer from one “existence” to another. For a given system of interest, they tell us how that system will evolve over time, in principle indefinitely. So the laws will naturally be an excellent authority on whether our constructed state leads to the evidence that we have actually observed.

“Okay,” you might think, “but why don’t we use *different* ampliative principles?” In other words, why do we use the *actual* laws to guide our reasoning here? After all, different laws could also link our constructed state with future states in which there are tire tracks of various lengths. That is true, of course, but using different ampliative principles would not be conducive

to our goal. We want to know what the *actual* speed of the car was. The *actual* laws are designed to be effective ampliative inference machines for actuality. Different ampliative principles would *not* accurately describe actuality. So to figure out whether the car was really traveling 30 miles per hour, we use the actual laws to see if the posited actual speed would have led to our actual evidence. It would not make sense to use ampliative inference machines that *don't* work for actuality to try to see whether a posited actual speed of 30 miles per hour would have led to our actual evidence.

Finally, why does our counterfactual reasoning use hindsight?⁶⁸ Suppose that we happen to know that some extremely unlikely chance event occurred that affected the length of the tire tracks. Suppose, for example, that the road was slick enough that the anti-lock braking system should have engaged, but that there was a very small chance that it would *not* engage, and in fact we have learned from the on-board computer that it did not engage. Also suppose that the chance of it not engaging is causally independent of how fast the car was moving.

Now suppose that our counterfactual reasoning did *not* use hindsight. We would then reason as follows. Construct a state similar to actuality in which the car is traveling 30 miles per hour. Evolve that state forward using the laws. The laws will generate two sets of models, one on which the anti-lock braking system engages, and one on which it does not. The length of the tire tracks will differ in these models; let the length in the first set be x , and the length in the second set be y .⁶⁹ The second set of models will be deemed extremely improbable, because the chance of the anti-lock braking system not engaging was very small. So we will then conclude that, if the

⁶⁸My argument here owes a great deal to Edgington (2004), which lays out a similar explanation for why our counterfactual reasoning uses hindsight.

⁶⁹For the sake of simplicity, I am supposing that all of the models in each set agree on the length of the tire tracks.

car was going 30 miles per hour, it was very probable that we would observe tire tracks of length x , and very improbable that we would observe tire tracks of length y .

But this is a misleading conclusion. Our goal is to figure out how fast the car was actually moving, and we know that in actuality, the anti-lock braking system did not engage. By not holding fixed the chancy non-engagement of the anti-lock braking system, we end up thinking that tire tracks of length y were less probable than they actually were (assuming the car was traveling 30 miles per hour). And since our aim is to reason backward from the observed length of the tire tracks to the speed of the car, we are therefore likely to reach an incorrect conclusion if we do not hold fixed the chancy non-engagement of the anti-lock braking system. The point of using hindsight, then, is to reach a more informed conclusion about the probability of the hypothesis of interest, given the evidence we have obtained.

Notice that I have said nothing yet about whether this constructed state, in which the car is traveling 30 miles per hour, is actual or counterfactual. Our goal in carrying out this reasoning process is precisely to *figure that out*. If it turns out that the observed tire tracks were extremely improbable on the supposition that the car was traveling 30 miles per hour, then we conclude that the car was not traveling 30 miles per hour, we deem the supposition *counterfactual*, and we infer that our constructed state is really an *altered* state. Alternatively, if it turns out that the observed tire tracks were extremely likely on the supposition that the car was traveling 30 miles per hour, then we conclude that the car was traveling at 30 miles per hour, we deem the supposition *factual*, and we infer that our constructed state is really a *past* state.

So reasoning in accord with the Altered States Recipe makes a great deal of sense if the purpose of such reasoning is to evaluate how strongly a given piece of unanticipated evidence bears on a given hypothesis. Of course, it should be mentioned that there are alternative ways we

can try to evaluate such evidential relations. For example, instead of constructing a past state in which the car was traveling 30 miles per hour and seeing whether it results in tire tracks of the observed length, we might instead plug the observed tire track length into the laws of motion and attempt to *retrodict* the speed of the car, i.e. run the laws backward to see what previous states must have preceded this present one.

This method of retrodiction might seem to have an advantage over evaluating counterfactuals using the Altered States Recipe. After all, if the laws are reverse-deterministic it tells us *precisely* which states must have preceded the present evidence. By contrast, even if the laws are deterministic, counterfactual reasoning delivers at best defeasible evidence. We can conclude that yes, if the car was traveling 30 miles per hour, it would have produced tire tracks of this length, but these tire tracks could also have been produced in other ways, e.g. if the road were slightly slicker and the car had been moving slightly slower.

But this difference is actually illusory. The problem is that retrodictions about states of the past based on information about the present are remarkably sensitive to very fine details about the present state. Put differently, for many physical processes, microscopic differences in the present state quickly ramify, via the laws, to macroscopic differences in past states. To see this, it can be helpful to imagine a video of a typical physical process played backward. For example, consider the dispersal of seeds from a dandelion seed head in the wind, and imagine viewing a video of this process playing out in reverse. Starting from a position in which the seeds are distributed across the grass downwind of the seed head, they all leap up from their landing spots and traverse winding paths through the air that eventually converge, miraculously, at their original positions in the seed head. And now imagine trying to use the position of the seeds in the grass downwind of the seed head to *retrodict*, using the laws, the original position of the seeds.

The problem is that to get even a *remotely* accurate retrodiction of the beginning state, we are going to require *extremely* detailed knowledge about the end state. For even a minute difference in the wind, say, or a small difference in the positions and energies of the molecules comprising the various seed heads and the blades of grass they rest on, will lead to an extremely different beginning state in which many of the seeds completely pass by the seed head rather than converging to their places in it, or in which they never leap up from the grass in the first place.⁷⁰

By contrast, using the laws to make *predictions* about future states from past states is usually *not* sensitive to fine details about the past state in question. To put it differently, for many physical processes, their evolution toward the future is largely unaffected by small changes to their initial conditions. For example, making a small change to the windspeed or the energies of the molecules comprising the seeds still results in an end state where the seeds are dispersed throughout the grass downwind of the seed head (albeit not in *exactly* the same way).⁷¹

Returning to our investigation of the speed of the car, the upshot of all this is that to retrodict the speed of the car to any degree of accuracy, we would have to know, not only the length of the tire tracks, but also the precise present state of variables like the molecules in the road, the tires, and maybe even the air around the wheels. In practice, we are not able to ascertain such facts, so even if we try to retrodict the speed of the car using deterministic laws, we will have to plug in a *set* of possible values for each of these variables, which will result in all sorts of different possible speeds of the car at the earlier state. But conversely, to use a posited state of the past to “predict” the length of the tire tracks to a reasonable degree of accuracy does not require this fine-grained knowledge of initial conditions.

⁷⁰Similar points are made, using different examples, by Loew (2017a), Elga (2000), and Horwich (1987).

⁷¹See Elga (2007) and Loew (2017a) for further discussion.

Lastly, reasoning in reverse is counterintuitive to us, and we are not particularly proficient at reaching verdicts about the plausibility of reversed temporal evolutions of physical systems. By contrast, counterfactual reasoning in accord with the Altered States Recipe is much more intuitive to us, because it requires us to reason from earlier states to later states, which is something that we do all the time. So there is no particular advantage—and indeed there are significant *disadvantages*—to using the laws to retrodict facts about actuality rather than using counterfactual reasoning to ascertain such facts. Given the choice, it is entirely reasonable that we usually opt for the latter method.

There are, of course, uses of counterfactual reasoning where we know ahead of time that the supposition is counterfactual. Two questions arise about these situations. First, why do we care about counterfactual situations when we already know they are counterfactual? And second, why do we still reason about these situations using the Altered States Recipe, especially considering that the stated motivations for using that recipe are absent?

One instance in which we might be interested in counterfactuals when we already know they are counterfactual is if we are trying to convince somebody *else* of some fact about actuality, given some evidence that they already accept. For example, maybe I already know that the car was traveling more than 30 miles per hour (maybe I was in the passenger seat), but my fellow investigators are not convinced by my report, because I suffered a concussion in the crash and my memory is now suspect. Thus I can go through the same reasoning as above to show them that if the car had been traveling at 30 miles per hour, we would have observed tire tracks of a different length. In this sort of situation, it still makes a great deal of sense to evaluate counterfactuals in accord with the Altered States Recipe.

Once this Recipe is in place—once counterfactual reasoning gets a foothold in our ways of thinking—we may be interested in evaluating all sorts of other counterfactuals. “What if there had been no moon? What if kangaroos had no tails? What if Oswald hadn’t shot Kennedy?” We know, of course, that these suppositions are genuine counterfactuals. So our interest in them can’t be based, in any straightforward way, on a desire to figure out facts about actuality. But I see no problem saying that we are interested in such counterfactuals simply because we are curious. There are practical reasons that we require counterfactual reasoning to begin with, but once the possibility of thinking in this way is available to us, it can be used in a variety of other contexts. Absent some alternative practical factors that would motivate reasoning about counterfactuals in a *different* manner, there is no particular reason to expect changes to the Altered States Recipe in cases where the original motivations are absent. It can be carried along intact simply by its own “conceptual inertia.”

4. Humean Ramifications

As I alluded to above, this explanation of counterfactual reasoning is particularly friendly to Humean accounts of laws, according to which the laws are regularities in the particular matters of fact. This may come as a surprise. The counterfactual resilience of the laws has traditionally been a challenge for Humean theories to explain. Roughly, the thought is that there has to be some unique feature of the laws that makes it appropriate for them to be held fixed in counterfactual reasoning. For example, on various *anti*-Humean theories (such as Maudlin’s, Armstrong’s, and Lange’s), the laws are held fixed because they have a special metaphysical status—a degree of *necessity*—that is not possessed by the particular matters of fact.⁷²

⁷²Strictly speaking, on Lange’s view the explanation runs in the other direction: the fact that the laws are held fixed in counterfactuals is what accounts for their necessity.

Conversely, if the laws are mere patterns in the particular matters of fact, then there is nothing metaphysically special about them—no degree of necessity—that could account for their characteristic degree of counterfactual resilience.

Of course, Humeans are free to just *stipulate* that, when we evaluate counterfactuals, the actual laws must be held fixed.⁷³ But without identifying some feature of the laws that makes them *suitable* to be held fixed in counterfactual reasoning, this sort of stipulation can feel pretty ad hoc.

However, the argument in §3 identifies a Humean-friendly feature of the laws that can explain their counterfactual resilience, namely, *extensive dynamical implications*. If the laws have extensive dynamical implications regarding systems in the actual world, and if our goal in counterfactual reasoning is to *figure out* facts about the actual world, then (as I have argued) it makes a great deal of sense for us to hold fixed the laws in counterfactual reasoning. It is crucial to note that nothing in this explanation requires the laws to be metaphysically privileged, for mere patterns in the particular matters of fact can have extensive dynamical implications (“states of type *X* are followed by states of type *Y*, which are followed by states of type *Z*”, etc.). So the argument in §3 is perfectly compatible with a Humean metaphysics.

Of course, that is not to say that it is compatible with just *any* Humean theory of laws. Rather, it requires a Humean theory that selects, as laws, patterns that have the appropriate dynamical form to figure into the Altered States Recipe. This arguably rules out a theory like Lewis’s Best System Account, according to which the laws are essentially just *efficient*

⁷³Lewis’s theory of counterfactuals is a prime example of this, though that theory only requires that the laws be held *mostly* fixed, as it allows for certain sorts of violations of actual law.

summaries of the particular matters of fact.⁷⁴ It is difficult to see why efficient summaries would have to have the sorts of dynamical implications that would allow us to calculate the temporal evolution of the altered state, and therefore it is doubtful that the laws of the BSA would be suited to play the role required by the Altered States Recipe.⁷⁵

But the argument in §3 does accord nicely with the Best Predictive System Account of laws, which mandates that the laws should have extensive dynamical implications and should apply, at least approximately, to isolated systems. Essentially, this explanation of counterfactual reasoning requires us to make “pseudo-predictions” about the systems of interest. Given a system S with observed property P_1 at a particular time, we want to know whether S also had some other property P_2 at an earlier time. To figure this out, we ask: “Would S have exhibited P_1 if it had exhibited P_2 ?” We then posit that S had P_2 at the earlier time, and “predict” whether such a system comes to have P_1 .

Finally, there is one other way in which the argument of §3 lends support to Humean theories of laws: it helps to explain why we may have anti-Humean intuitions. Anti-Humeans are fond of pointing out that the Humean view of laws is at odds with our pre-theoretic intuitions. After all, it is hard to avoid the feeling that the laws are somehow *distinct* from the patterns in the particular matters of fact, and that the patterns could remain the same while the laws varied.

One common Humean response is to suggest that their concept of law is slightly revisionary, but that Humean laws do a good enough job playing the variety of roles attributed to

⁷⁴See, e.g., Lewis (1986), (1994). There are also a variety of developments of Lewis’s view that leave the basic picture of laws as efficient summaries intact, e.g. Loewer (1996), Cohen and Callender (2009), and Demarest (2017).

⁷⁵This is essentially the same argument that I gave in Chapter 2, to the effect that the laws of the BSA are not likely to be predictively useful to creatures like us.

the laws in scientific practice that they are still deserving of the name.⁷⁶ I think this is too concessive: the admission of revisionism should be avoided if at all possible. And in this case, I think Humeans can avoid it. Particularly, we can use the Best Predictive System Account of laws to help explain why we have anti-Humean intuitions in the first place.

I have little doubt that there are a variety of sources to our anti-Humean intuitions, but one particularly powerful source seems to be the laws' counterfactual resilience. The laws would still have been true under a very wide range of counterfactual suppositions. And since these counterfactual changes to worldly phenomena do not engender any violations of actual law, it is only natural to think that the laws must stand "above" or "behind" those phenomena. In other words, the counterfactual resilience of the laws encourages us to view them as unaffected by the various goings-on in the physical universe, and therefore to think of them in an anti-Humean light.

But if the foregoing arguments are correct, the Humean has a way of explaining the laws' counterfactual resilience in a way that is friendly to her own view of laws. The laws are counterfactually resilient, not because of any special metaphysical status they possess, but because their predictive utility makes it useful for us to hold them fixed in order to achieve our goal in counterfactual reasoning. The Humean is therefore able to account for a significant source of our anti-Humean intuitions, even on the assumption that Humeanism is correct. Thus, while it is tempting to think that violations of our intuitions constitute a serious theoretical cost for Humeanism, the account of counterfactual reasoning that I am giving here (coupled with the Best Predictive System Account of laws) drastically reduces that cost.

⁷⁶See, for example, Hall (ms, p. 25). Barry Loewer has also expressed this thought to me in personal conversation.

5. Conclusion: Does God Think About Counterfactuals?

Philosophers have often puzzled about why creatures like us should care about counterfactuals. Shouldn't we just care about actuality? Suppose we knew all the facts about actuality: what utility could there possibly be to learning facts about non-actuality? The worry is exacerbated by thinking about counterfactuals in terms of possible worlds (which has become standard since Lewis [1973]): why should thinking about \textit{other possible worlds} be of any help in navigating *this* one?

My arguments here provide an answer to that question. When we engage in counterfactual reasoning, we often do not *know* that we are doing so. It is only in hindsight, after we have evaluated the counterfactual and concluded that our observed evidence would not have occurred on the given supposition, that we conclude that we were evaluating a counterfactual all along. Furthermore, the goal of such reasoning is to ascertain facts about the actual world. So if we *did* know all the facts about the actual world, there would be no point in engaging in counterfactual reasoning. To put it pithily, God does not think about counterfactuals. Only we do.

CHAPTER 5: CONCLUSION

By way of concluding, I want to first suggest some further avenues for developing and applying the Best Predictive System Account, and then I will wrap things up with some all-too-grandiose remarks about two different understandings of scientific theorizing.

Back in Chapter 1, I listed a number of conceptual roles played by the laws of nature in scientific practice. Here is that list again:

1. Laws support counterfactuals
2. Laws have a characteristics necessity, but they are also contingent in some respect
3. Laws are important features of our world worth knowing
4. Laws are things that scientists try to discover using the relevant epistemic principles
(whatever those may be)
5. Laws are useful for predictive purposes
6. Laws underwrite reliable measurement procedures
7. Laws explain natural phenomena
8. Laws are confirmed by their instances
9. Laws govern the evolution of events in the universe

I argued, back in Chapter 1, that any philosophical theory of laws ought to proceed by first identifying some natural *subset* of these roles to be constitutive of lawhood. And at that point the task is to explain why, given that core conception of laws, they would naturally play those other *non-constitutive* roles as well. The success of a theory is then a function, both of how many of

these additional roles can be explained by that core conception, and by how *convincing* or *illuminating* or *non-ad-hoc* those explanations are.

Let me now use this theoretical measuring stick to recap where the Best Predictive System Account stands. I have argued that we should take roles (3), (4), and (5) as definitive of our core conception of laws. The laws are, first and foremost, a set of principles that is maximally predictively useful to creatures like us. Their predictive utility (role 5) makes them worth knowing (role 3), and as we saw in Chapter 2, it also explains why scientists use the particular standards that they do when they are investigating the laws (role 4). In Chapter 3, then, I argued that *given* this core conception of laws, we can give a natural explanation for why the laws support counterfactuals (role 1). Along the way, we also gained a deflationary explanation of why we may come to associate anti-Humean intuitions about *governance* with the laws (role 9). Furthermore, though I have not yet made this argument, we may be able to use the counterfactual resilience of the laws to explain their characteristic degree of *necessity* (role 2), much along the lines of Lange's (2009) discussion (albeit proceeding from a very different metaphysical starting point).

Here, then, are the remaining roles on our list. Role 6: the laws underwrite the reliability of measurement procedures; role 7: the laws explain natural phenomena; and role 8: the laws are confirmed by their instances. None of these, I submit, is a straightforward consequence of the Best Predictive System Account. So here I have plenty of work cut out for me. In each case, my task is to explain why the laws would play these roles, if the laws are what the BPSA says they are. I do not have time to discuss all three of these roles here, but let me gesture at how the explanation *might* go in the case of role 7.

The laws of nature are supposed to explain natural phenomena—in particular, they are supposed to be able to explain their *instances*. For example, the Lorentz force law states that a charged particle traversing a magnetic field will experience a Lorentz force perpendicular to its direction of motion. Thus, if the laws play the explanatory role that we traditionally attribute to them, we should be able to explain why any particular charged particle experiences a Lorentz force when it enters a magnetic field by appealing to the Lorentz force law.

Now, there has recently been an explosion of interest in whether *reductive* laws are suited to play this explanatory role. Roughly, the worry is that since reductive laws are just patterns in the particular matters of fact, it's not clear that they can also *explain* the particular matters of fact, because that looks dangerously close to self-explanation.⁷⁷ So one challenge for a reductive account of laws like my own is to find some way to alleviate this circularity concern.

But it would also be prudent to take a broader perspective, and try to figure out why creatures like us *care* about explanatory considerations in the first place. For example, when two theories both appear to be empirically adequate, we prefer the more explanatory one to the less explanatory one. And it seems worth considering why that is the case.

One question right off the bat is: What features of a theory make it count as more explanatory? Part of the project here would involve trying to answer this question—that is, trying to get a clearer sense of what features lend a theory more or less explanatory power. Once we get some clarity on this point, we may get a better idea of why creatures like us prefer theories that are more explanatory, other things being equal.

In that vein, suppose we are comparing two theories, Theory *A* and Theory *B*. And suppose that, as far as we know both theories are empirically adequate. That is, both of them are

⁷⁷For more detailed statements of the worry, see, for example, Lange (2013, 2018), Hicks and van Elswyk (2015), Marshall (2015), Miller (2015), Shumener (forthcoming) and Roberts (ms).

able to “cover” or “subsume” all of the phenomena that we have encountered so far. But suppose that the *manner* in which each theory covers these phenomena is very different. On the one hand, Theory *A* covers a wide range of phenomena just by appealing to miraculous, fine-tuned arrangements in the initial conditions of the universe. In other words, when we ask Theory *A* to explain why a given phenomenon occurred, it typically ends up saying that the phenomenon occurred because the initial conditions of the universe were just right. By contrast, suppose that Theory *B* covers the same range of phenomena, but it does so largely (or purely) by appealing to *dynamical* considerations. In other words, it says that no matter what the initial conditions of the universe were, the dynamics make the occurrence of these phenomena very likely. In such a case, we will prefer Theory *B* to Theory *A* on the grounds that Theory *B* gives better explanations of the same range of phenomena. But *why* are Theory *B*’s dynamical explanations better than Theory *A*’s initial condition explanations?

One thought is that Theory *B* renders the phenomena to be explained much more probable, whereas Theory *A* renders those same phenomena *improbable*, because the initial conditions had to be fine-tuned *just so* in order for those phenomena to occur. So maybe our preference for Theory *B* here is just a matter of Theory *B* rendering our observed evidence more probable than Theory *A*.

I don’t want to suggest that there is anything wrong with this account of our preference for Theory *B*. But another way of accounting for that preference is instead to appeal to Theory *B*’s greater *predictive utility*. More specifically, suppose we are hoping to use either Theory *A* or Theory *B* to generate predictions about future phenomena. In that case, Theory *B* is going to lend itself much more readily to making these sorts of predictions. Since its explanations rely purely (or largely) on dynamical considerations, it should not be too difficult to extract further

predictions out of it. By contrast, Theory *A*'s explanations primarily rely on minute details about the universe's initial conditions, about which we are unlikely to have very detailed knowledge. So it would appear to be quite difficult to extract further predictions out of Theory *A*.

This, then, is the general hope. Once we get clear about what sorts of features make a theory more explanatory, we will see that those features are actually roundabout ways of making the theory more predictively useful. Obviously, this idea accords very nicely with the Best Predictive System Account of laws. But just as obviously, it needs much more development before a convincing argument can be mounted on its behalf.

Before wrapping things up, I want to suggest one further application of the Best Predictive System Account. If the BPSA is correct, it provides us with a counterargument to a prominent theory in the philosophy of statistical mechanics.

David Albert and Barry Loewer have recently been developing an extraordinarily ambitious project that aims to explain the direction of time purely in terms of the fundamental laws of nature.⁷⁸ The project is based on a collection of principles that Loewer calls "the Mentaculus." The Mentaculus consists of (i) the fundamental dynamical laws of nature, (ii) the Past Hypothesis (PH), and (iii) the Statistical Postulate (SP). Together, PH and SP say that the early universe began in a very low entropy macrostate, and that there is a uniform probability distribution over the possible microstates that could have realized that macrostate. In other words, they amount to a set of probabilistic constraints on the initial conditions of the universe. According to Albert and Loewer, we should treat PH and SP as additional, non-dynamical laws of nature. That is, Albert and Loewer think that it is a matter of law that certain initial conditions of the universe were more probable than others.

⁷⁸See, for example, Albert (2000, 2015), and Loewer (2001, 2004, 2012).

Albert and Loewer hope to use the Mentaculus to explain the direction of time. They think the direction of time actually consists of a *variety* of distinct temporal asymmetries in our everyday experience of the world, such as the asymmetries of knowledge and control. The asymmetry of knowledge consists in the fact that we have much more detailed knowledge about the past than about the future. And the asymmetry of control consists in the fact that we seem to be able to control facts about the future, but not about the past. The thought is that if we can explain these temporal asymmetries, we will have explained the direction of time. And we can explain these asymmetries, they think, by appeal to the Mentaculus.

The details of this explanation are not fully developed, but here is a rough sketch.⁷⁹ If we follow Albert and Loewer's recommendation to regard all of the elements of the Mentaculus as laws of nature, then those elements will be held fixed in counterfactual reasoning. Consequently, counterfactual suppositions will tend to produce very minor differences moving toward the past, because any such differences will be constrained by the lawhood of PH and SP, which act as constraints on the universe's initial conditions. By contrast, there is no analogous constraint on the \textit{future} conditions of the universe—there is no future analog of PH and SP. Rather, the only laws constraining future states are the fundamental dynamical ones. So counterfactual differences in past states are likely to produce profound differences in future states, but counterfactual differences in future states are *unlikely* to produce any significant differences in past states (especially in past *macroscopic* states).

In short, the Mentaculus, coupled with the fact that the laws of nature are held fixed in counterfactual reasoning, results in an asymmetry of counterfactual dependence: the future

⁷⁹In addition to the above sources, see Loew (2017b) for a more detailed discussion.

counterfactually depends on the past, but not vice versa.⁸⁰ Albert and Loewer's hope, then, is to use this counterfactual asymmetry to ground the various more specific temporal asymmetries, such as the asymmetries of knowledge and control.

I do not think this explanation can work. The general worry is this. As I have argued in the preceding chapters, the most promising way to reach a philosophical understanding of laws is to take our epistemic situation as given, and develop an account of laws informed by that situation. In particular, I have argued that the laws are principles that creatures like us would find maximally predictively useful. But what makes something predictively useful to us depends crucially on certain manifest facts about our epistemic situation. For example, one very prominent fact about our epistemic situation is that we have far more detailed knowledge about the past and present than we do about the future. Indeed, this lack of knowledge about the future is part of what makes us interested in discovering the laws in the first place! So the project of appealing to the laws of nature to explain the direction of time gets things *backward*. We have to *assume* the direction of time, and the various more specific temporal asymmetries, to develop a philosophical understanding of the laws of nature.

This general worry can be put in a more specific form. One of the crucial moves in Albert and Loewer's explanation of temporal asymmetries is the claim that the laws of nature are held fixed in counterfactual reasoning. But their discussion ignores the question of *why* the laws of nature are held fixed in counterfactual reasoning. In Chapter 3, I argued that the laws are held fixed because doing so is useful to creatures in our epistemic situation. Thus, if we supplement Albert and Loewer's explanation of temporal asymmetries with my account of why the laws are held fixed in counterfactual reasoning, here is how it would go:

⁸⁰More carefully, the past does not counterfactually depend on the future to an appreciable degree

1. All of the elements of the Mentaculus are laws of nature.
2. The laws of nature are held fixed in counterfactual reasoning because doing is useful to creatures in our epistemic situation.
3. Holding the elements of the Mentaculus fixed in counterfactual reasoning results in an asymmetry of counterfactual dependence.
4. The asymmetry of counterfactual dependence explains why we're in the epistemic situation that we're in (namely, why there are temporal asymmetries of knowledge, control, etc.)

The problem here is that this argument only explains why we are in the epistemic situation that we're in by first *assuming* (in Premise 2) that we are in that very situation. And this strikes me as deeply unsatisfactory.

Hopefully these project sketches illustrate the fruitfulness of the BPSA. Now to conclude, I want to briefly consider two different understandings of scientific theorizing.

It is widely acknowledged that the laws of nature have important conceptual connections to a host of other notions such as counterfactuals, chances, dispositions, and so on. Following Briggs and Forbes (2018), let us call the collection of all these notions the “nomological package.” Scientists, clearly, aim to discover aspects of the nomological package. They investigate the laws of nature, evaluate various counterfactuals, discover the chances of various phenomena, and so on. But *why* do scientists aim to discover the nomological package?

One answer is just that scientists want to know as much as they can about the world, and that *includes* truths about the nomological package. In addition to all the particular matters of occurrent fact, the laws, counterfactuals, chances, dispositions, and so on are all just a *part* of the totality of facts about the world, so in their fervent thirst for knowledge, scientists naturally want

to know about the elements of the nomological package. On this view, laws, counterfactuals, etc. are all perfectly objective: any observer, whether human, god, or oracle, would have equal reason to be interested in them. Call this the *objective conception* of the nomological package.

A different answer is that scientists *primarily* care about the particular matters of occurrent fact. But it turns out that creatures like us, if left to our own perceptual capacities, don't have very effective ways of *coming to know* much about the particular matters of occurrent fact. Thus we need to invent *tools* to help us do so. That, then, is the task of the elements of the nomological package: they are instruments to help increase our knowledge of the particular facts. Call this the *ampliative conception* of the nomological package.⁸¹

On the ampliative conception, the nomological package is not guaranteed to be inherently interesting to every observer. For example, gods and oracles have no need for it, since it is a tool designed to help us see further, and they already know everything that it could ever reveal to us. But of course, the nomological package *is* interesting to *us*, because it consists of exactly the sorts of theoretical tools that we need to figure out more about the *real* subject of interest: the particular matters of occurrent fact.

Many of my arguments in the preceding chapters have tried to point out the advantages of adopting the ampliative conception. But really what I want to argue for is its *tenability*. I think the ampliative conception of the nomological package has the potential to provide some particularly illuminating and fascinating understandings of scientific practice, and so far this potential has been relatively underexplored. Of course, time will tell if it stands up to scrutiny. But I think its prospects are promising.

⁸¹ Adopting the ampliative conception need not imply that claims about laws, counterfactuals, etc. don't have truth values, or that they are purely instrumental. Rather, it may be that their truth is simply a *function* of their utility in ampliative inferences.

APPENDIX

In this appendix, I address several additional objections to my characterization of the Humean base as the totality of macroscopic phenomena. Many of these objections will be familiar, as I am largely drawing them from the literature on van Fraassen's constructive empiricism. van Fraassen (1980) famously claimed that the aim of science was to generate theories that are empirically adequate with respect to the observable phenomena. Thus, some of the objections to van Fraassen's theory can be adapted against my own. While I do not have room to address all of the potential worries here, I will address what I take to be some of the more serious and interesting ones.

First, there are a variety of objections that focus on constructive empiricism's agnosticism about unobservable entities. Here I include Putnam's miracle argument, namely the claim that scientific realism about unobservables is "the only philosophy that doesn't make the success of science a miracle" (1975, p. 73). In a very similar vein, supporters of inference to the best explanation would argue that the best explanation for the success of science is that there really are the unobservable entities that it posits. And with these I will group one of Maxwell's (1962) arguments: even granting a principled distinction between the observable and the unobservable (which he actually refuses to grant), this has nothing to do with *existence*. It may seem that all of these arguments can be modified to challenge my view if we simply swap out "unobservable entities" for "microscopic phenomena."

However, none of these arguments actually applies to my view, even in their modified forms. All of them attack constructive empiricism on the grounds that it refrains from

committing to the existence of unobservable entities. My view does not refrain from such commitment. This was partly why I was reluctant to phrase my view using Hall's terminology of

“manufactured” entities that don’t “really” exist. On my view, protons, electrons, quarks, etc. *really do* exist. It is only that the criteria for them “really existing”—what makes it true that they really exist—are more complex than we might have naïvely thought they were.

Nevertheless, there is a sense in which the spirit of these objections still does apply to my view. After all, while I do maintain that, for example, there really was a proton that produced track (6) in Figure 2, I also maintain that the *truthmaker* for that claim is not simply that there was some microscopic particle with all the properties of a proton in the bubble chamber. The truthmaker, rather, is that the existence of protons falls out of the Nomic Formula when it is applied to the totality of macroscopic phenomena. As I put it earlier: it is because a variety of phenomena (including the Brookhaven bubble chamber results, but also including a plethora of other phenomena from other experiments) behave so much as if there were microscopic particles with the properties that we ascribe to protons that our claims about protons come out true.

So perhaps the objection can be re-framed to address my theory, as follows. My theory holds that the truthmakers for claims about microscopic phenomena are not nearly as straightforward as the truthmakers for claims about macroscopic phenomena. Yet this alleged semantic difference is not marked by any significant syntactic difference in our language; the sentence “there is a proton” is syntactically quite similar to the sentence “there is a chair.” And other things being equal, we should posit similar semantic structures where we find similar syntactic structures. Thus, if other things are equal, my theory posits a gratuitous semantic difference where there is no syntactic difference. Furthermore, *things really are equal*. Both Putnam’s miracle argument and inference to the best explanation demonstrate that we have good reason to think that there really are microscopic particles that can serve as the referents of our term “proton,” just as there really are chairs that can serve as the referents of our term “chair.”

I am no fan of the dictum to posit similar semantic structure where there is similar syntactic structure. Tacit adherence to this dictum helps generate what Gilbert Ryle variously called “category mistakes” (2000) and “wrongly imputed parities of reasoning” (2015). Recall his example (2000, p. 16) of a foreigner visiting Oxford and getting a tour of the colleges, libraries, laboratories, etc. After all this, the foreigner confusedly asks to be shown the university. To use Ryle’s terminology, the foreigner has wrongly allocated the concept of a university to the same logical category as that of libraries, laboratories, and so on. Now, no one doubts that the foreigner really is making a mistake here—a mistake whose basis is a linguistic misunderstanding. Yet the dictum to posit similar semantic structure where there is similar syntactic structure would encourage exactly this misunderstanding: “university” is a noun just like “library” and “laboratory,” so the university must be another facility like the libraries and the laboratories. Therefore, I maintain, this dictum should be disregarded. It is perhaps an acceptable, defeasible guideline for someone learning a new language, but it is out of place as a regulative principle in philosophical theorizing.

Another objection relates to my previous claim that my theory avoids the possibility that the laws are underdetermined by all possible evidence, a worry that afflicts both the Lewisian BSA and non-reductionist theories of laws. This objection maintains that my theory countenances a very similar sort of possibility. In particular, some macroscopic phenomena are so remote in space and time that they are epistemically inaccessible to us. Yet according to my theory, those phenomena are supposed to be plugged into the Nomic Formula to yield the laws. So according to my theory, to figure out the laws we must apply the Nomic Formula to facts that we could never apply it to. Why is this not problematic for my theory in the same way that, for example, Lewis’s theory is problematic in allowing for in-principle-undetectable fundamentalia?

This is a challenging objection, though it is not exactly a new one. Supporters of the BSA have long touted its epistemological benefits over non-reductionist theories, and it has long been open to non-reductionists to respond along the lines I have just suggested. Tim Maudlin is about as non-reductionist as they come, and he articulates exactly this kind of response:

Let us first recognize that the existence of facts which are not determined by the complete totality of all observations, past, present, and future, is commonplace. No doubt, Socrates had a blood type. Also, no doubt, one could not deduce that blood type from a complete catalogue of every observation that has been or will be made[...] Socrates' blood type is now, and will always be, beyond our ken. Of course, only a lunatic would conclude that the ontological status of Socrates' blood is thereby affected, that he (miraculously) had no blood type at all, or an indeterminate one. No doubt there was a fact of the matter about his blood, despite our irremediable inability to know it. So the question arises: what difference does it make to ontology that a particular fact cannot be deduced from actual evidence? (2007, pp. 73-74)

As Maudlin suggests, we needn't go all the way out to extremely remote facts in order to generate this kind of worry. Indeed, consider just the set of all phenomena that we ever *actually observe*. There is a sense—a very restricted modality—in which any other phenomena are epistemically inaccessible. For we would have to be in a possible world different from the one we are actually in to observe them. Thus, if it is a legitimate desideratum in constructing a philosophical theory of laws to minimize the possibility of underdetermination, why not just claim that the Humean base consists of the totality of *actually observed* phenomena? This would preclude the possibility of *any* in-principle-undetectable errors.

In short, the response is this: minimizing the possibility of underdetermination is *one* desideratum in constructing a philosophical theory of laws. But there are others as well, and it turns out that the account I just elaborated would strongly conflict with them. The problem is that it does not make sense of actual scientific practice. And making sense of scientific practice is at least as weighty a desideratum as minimizing the possibility of underdetermination.

Why does this account, according to which the Humean base is the totality of actually observed phenomena, not make sense of scientific practice? According to this account, to figure out what the laws of nature are, there would be no need to make any new observations or perform new experiments. We could simply stop looking for new data and apply the Nomic Formula to the totality of actually observed phenomena. Indeed, if this account were correct, we might see scientists attempting to *prevent* new experiments in order to preserve the currently-accepted theories. And doing so would not be considered *bad science*. It would simply be a way of trying to ensure that the currently-accepted theories are correct. But scientists manifestly *do* go out of their way to gather new observations and perform new experiments in order to investigate the laws of nature. Preventing new experiments manifestly *is* bad science (if it can be called “science” at all). This is what I mean by saying that this account does not make sense of scientific practice.

The challenge, then, is to create an account that makes sense of scientific practice while minimizing the possible sources of underdetermination. I maintain that my account accomplishes this. On the one hand, the underdetermination countenanced by my account is not gratuitous. Consider, for example, macroscopic phenomena occurring in the most remote galaxy in the universe. Our current theories tell us that we have no way of ever ascertaining those facts, but these theories could be wrong (however unlikely that currently seems). And if it somehow turns out that we can access such facts, they would certainly constitute relevant data points in our theory building. My theory also makes sense of why scientists attempt to perform new experiments and gather new data: that data constitutes further facts about macroscopic phenomena, which are relevant to figuring out the laws of nature.

On the other hand, my account introduces fewer sources of underdetermination than Lewis's. The BSA introduces too many possible sources of underdetermination by making the Humean base consist of the totality of facts about fundamentalia. Granted, there is still a possibility of in-principle-undetectable errors on my account. But the set of such possible errors countenanced by my account is a proper subset of the set of such possible errors countenanced by Lewis's account. That is, any in-principle-undetectable fact about macroscopic phenomena is a potential source of error about the laws on my account, and it is likewise a potential source of error on Lewis's account. After all, macroscopic phenomena are comprised of facts about fundamentalia. But Lewis also countenances further possible errors; assuming we have gathered *all* of the facts about macroscopic phenomena, we still might not have all the facts about the fundamentalia, and thus we could still be wrong about the laws, according to his view. Unlike the errors allowed by my account, these errors are not needed to make sense of scientific practice.

The last objection I want to consider is based on Maxwell's (1962) "pseudohistorical" recap of the germ theory of disease. According to Maxwell's pseudohistory, a "Pasteur-like" fellow named Jones had the idea that diseases were spread by minuscule entities called "crobes," which were too small to be seen. On the basis of this theory, he derived a number of preventative measures, which, upon their eventual adoption, were extremely successful at reducing the number of contracted diseases. Philosophers of the time adopted a variety of the standard positions about unobservable entities with respect to Jones's crobes: realism, anti-realism, instrumentalism, etc. But within Jones's lifetime, the compound microscope was invented, and the microbes (as they came to be called) were observed in exquisite detail. Philosophers then had a variety of reactions. Some adopted increasingly idealist positions, maintaining that all

meaningful descriptive language was expressible in terms of sense data. Others argued that the microbes actually were never observed, since they had to be viewed through a microscope and not with the naked eye. Some others, says Maxwell, “freely admitted error and were converted to realist positions” (p. 187).

What error were they admitting to? Well, it depends of course on what their original position had been. Those anti-realists who had denied the existence of microbes certainly had a straightforward error to admit to, as did the instrumentalists, for here were the microbes, staring them in the face. Now the pertinent question for me is this: suppose I had advanced my theory before the invention of the compound microscope. Would I likewise have had to admit error once we looked through it and saw the microbes?

It may seem that I would. After all, I have maintained that the truthmakers for claims about microscopic phenomena are not facts about the microscopic entities themselves. According to my theory, claims about microscopic phenomena are made true because they are consequences of the best systematization of the *macroscopic* phenomena. This would seem to be in conflict with the discovery of such things as microbes, which could serve as the truthmakers for Jones’s claims about microbes. Furthermore, this would seem to be a unique problem for my theory as opposed to Lewis’s, since Lewis did not demur from a straightforward account of the truths about microscopic phenomena in the way that I have.

However, the visualization of microbes through the compound microscope poses no special challenge for my theory. Our observations through the compound microscope merely get grouped with the rest of the macroscopic phenomena. All they do is make it exceedingly likely that the Nomic Formula will output claims about microbes. That is, these observations make it a near certainty that microbes will figure amongst the propositions of the best system. The

truthmakers for claims about microbes are not the microscopic entities themselves; such claims are made true by being outputs of the Nomic Formula when it is applied to the totality of macroscopic phenomena. It merely turns out that some of the macroscopic phenomena (namely, what we see when we look through a compound microscope) provide us with a particularly good reason for saying there are such things as microbes.

It is also worth considering how we decided that the compound microscope was a reliable way to visualize microscopic phenomena. Obviously the device was not invented by chance; a great deal of optical theory was involved in its construction, on the basis of which it was *expected* to help us see things that were otherwise too small to be seen. No doubt the device also had to be tuned and adjusted in various ways in order to see the crobes. But suppose we had looked through the first compound microscopes and observed nothing that could correspond to Jones's crobes. We would then have to choose whether to abandon realism about crobes, or some element of optical theory, or, most likely, the belief that the microscope was constructed and tuned correctly. In other words, I am suggesting that it may have been partly *because* we already accepted Jones's crobe theory that we came to regard the compound microscope, adjusted in exactly these ways, as a reliable method of visualizing microscopic entities.

This point is perhaps easier to appreciate if one considers scanning tunneling microscopes (STMs) and the images of atoms that they produce. Before the invention of STMs in the 1980s, we had already long accepted the Democritean picture of atoms as something like minuscule billiard balls. Now, how did we decide that the images produced by STMs were correct? Well, partly because of the background quantum mechanical theory used to construct them. And partly, no doubt, because the images they produced agreed so readily with our previous beliefs about what atoms "look like." If STMs had produced images of atoms that looked like dashed lines

instead of spheres, we probably wouldn't have concluded, "Oh, atoms are apparently dashed line segments." No, we would first have attributed the error to improper tuning of the measuring device. If it somehow proved impossible to tune the device in a way to get plausible images, only then would we have had to choose between (a) the elements of the quantum mechanical theory used to construct it, and (b) our firmly-held atomic picture.

My goal in bringing up this last point is to emphasize that the relationship between novel imaging devices and previously-held theories about unobservables is not as straightforward as it may appear. It is tempting to view the observation of microbes through the compound microscope as a direct refutation of any theory that, like mine, rejects a standard account of the truthmakers for claims about microbes. I do not think the situation is that simple. There is a subtle yet significant difference between the following goals:

1. Construct a device to figure out whether there really are crobes
2. Construct a device that allows us to visualize crobes

My suggestion is that the construction of new observation devices is often more along the lines of (2) than it is (1).

REFERENCES

- Albert, David. (2000) *Time and Chance*. Cambridge, Massachusetts: Harvard University Press.
- Albert, David. (2015) *After Physics*. Cambridge, Massachusetts: Harvard University Press.
- Albert, David, and Barry Loewer. (1988) "Interpreting the Many Worlds Interpretation." *Synthese*. 77: 195-213
- Armstrong, David. (1983) *What is a Law of Nature?* Cambridge: Cambridge University Press.
- Aspect, Alain, Phillippe Grainger, and Gérard Roger. (1981) "Experimental Tests of Realistic Local Theories via Bell's Theorem." *Physical Review Letters*. 47: 460-3.
- Barnes, V. E., Connolly, P. L., et. al. (1964) "Observation of a Hyperon with Strangeness Minus Three." *Physical Review Letters*. 12: 204-6.
- Beebe, Helen. (2000) "The Non-Governing Conception of Laws of Nature." *Philosophy and Phenomenological Research*. 61: 571-94.
- Bhagal, Harjit, and Perry, Zee. (2017) "What the Humean Should Say About Entanglement." *Noûs*. 51: 74-94.
- Blackburn, Simon. (1993) *Essays in Quasi-Realism*. Oxford: Oxford University Press.
- Bohm, David. (1957) *Causality and Chance in Modern Physics*. Philadelphia: University of Pennsylvania Press.
- Briggs, R. A., and G. A. Forbes (2018) "The Future, and What Might Have Been." *Philosophical Studies*. <https://doi.org/10.1007/s11098-017-1026-y>
- Cartwright, Nancy. (1999) *The Dappled World: A Study of the Boundaries of Science*. Cambridge: Cambridge University Press.
- Chew, Geoffrey. (1962) "S-Matrix Theory of Strong Interactions without Elementary Particles." *Reviews of Modern Physics*. 34: 394-400.
- Chisholm, Roderick. (1955) "Law Statements and Counterfactual Inference." *Analysis*. 15: 97-105.
- Cohen, Jonathan, and Callendar, Craig. (2009) "A Better Best System Account of Lawhood." *Philosophical Studies*. 145: 1-34.
- Demarest, Heather. (2017) "Powerful Properties, Powerless Laws." In Jacobs, Jonathan (ed.), *Putting Powers to Work: Causal Powers in Contemporary Metaphysics*. Oxford: Oxford University Press.

- Earman, John, and Roberts, John. (2005) "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: 253-86.
- Edgington, Dorothy. (2004) "Counterfactuals and the Benefit of Hindsight." In Phil Dowe and Paul Noordhof (eds.), *Cause and Chance: Causation in an Indeterministic World*. New York: Routledge.
- Einstein, Albert. (1948) "Quanten-Mechanik Und Wirklichkeit." *Dialectica*. 2: 320-4.
- Elga, Adam. (2000) "Statistical Mechanics and the Asymmetry of Counterfactual Dependence." *Philosophy of Science*. 68: 313-24.
- Elga, Adam. (2007) "Isolation and Folk Physics." In Huw Price and Richard Corry (eds.), *Causation, Physics, and the Constitution of Reality*, pp. 106-19. Oxford: Oxford University Press.
- Eve, Arthur. (1939) *Rutherford: Being the Life and Letters of the Right Hon. Lord Rutherford, O. M.* Cambridge: Cambridge University Press.
- Feynman, Richard. (1965) *The Character of Physical Law*. New York: Modern Library.
- Fowler, William and Nicholas Samios. (1964) "The Omega-Minus Experiment." *Scientific American*. 211: 36-45.
- Geiger, Hans and Ernest Marsden. (1913) "The Laws of Deflexion of α Particles through Large Angles." *Philosophical Magazine*. 25: 604-23.
- Gell-Mann, Murray. (1964) "A Schematic Model of Baryons and Mesons." *Physics Letters*. 8: 214-5.
- Glymour, Clark. (1980) *Theory and Evidence*. Princeton, NJ: Princeton University Press.
- Goodman, Nelson. (1983) *Fact, Fiction, and Forecast*. Cambridge, MA: Harvard University Press.
- Hall, Ned. (ms) "Humean Reductionism About Laws of Nature."
<https://philpapers.org/archive/HALHRA.pdf>
- Hicks, Michael and Peter van Elswyk. (2015) "Humean Laws and Circular Explanation." *Philosophical Studies*. 172: 433-43.
- Horwich, Paul. (1987) *Asymmetries in Time*. Cambridge, MA: MIT Press.

- Hume, David. (1993) *An Enquiry Concerning Human Understanding*.} Ed. Steinberg, Eric. Indianapolis: Hackett.
- Jackson, Frank. (1977) "A Causal Theory of Counterfactuals." *Australasian Journal of Philosophy*. 55: 3-21.
- Knapp, Alex. (2012) "How Much Does it Cost to Find a Higgs Boson?" *Forbes*. Online Article. <https://www.forbes.com/sites/alexknapp/2012/07/05/how-much-does-it-cost-to-find-a-higgs-boson/#4f5d83453948>
- Ladyman, James, and Don Ross. (2007) *Everything Must Go: Metaphysics Naturalized*. New York: Oxford University Press.
- Lange, Marc. (2000) *Natural Laws in Scientific Practice*. Oxford: Oxford University Press.
- Lange, Marc. (2007) "Laws and Meta-Laws of Nature: Conservation Laws and Symmetry." *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*. 38: 457-81.
- Lange, Marc. (2009) *Laws and Lawmakers*. New York: Oxford University Press.
- Lange, Marc. (2013) "Grounding, Scientific Explanation, and Humean Laws." *Philosophical Studies*. 165: 255-61.
- Lange, Marc. (2018) "Transitivity, Self-Explanation, and the Explanatory Circularity Argument against Humean Accounts of Natural Law." *Synthese*. 195: 1337-53.
- Lewis, David. (1973) *Counterfactuals*. Malden: Blackwell.
- Lewis, David. (1979) "Counterfactual Dependence and Time's Arrow." *Noûs*. 13: 455-76.
- Lewis, David. (1983) "New Work for a Theory of Universals." *Australasian Journal of Philosophy*. 61: 343-77.
- Lewis, David. (1986) *Philosophical Papers, Volume II*. New York: Oxford University Press.
- Lewis, David. (1994) "Humean Supervenience Debugged." *Mind*. 103: 473-490.
- Loew, Christian. (2017a) "Causation, Physics, and Fit." *Synthese*. 194: 1945-65.
- Loew, Christian. (2017b) "The Asymmetry of Counterfactual Dependence." *Philosophy of Science*. 84: 436-55.
- Loewer, Barry. (1996) "Humean Supervenience." *Philosophical Topics*. 24: 101-27.

- Loewer, Barry. (2001) "Determinism and Chance." *Studies in the History and Philosophy of Modern Physics*. 32: 609-29.
- Loewer, Barry. (2004) "David Lewis's Humean Theory of Objective Chance." *Philosophy of Science*. 71: 1115-25.
- Loewer, Barry. (2007) "Laws and Natural Properties." *Philosophical Topics*. 35: 313-28.
- Loewer, Barry. (2012) "Two Accounts of Laws and Time." *Philosophical Studies*. 160: 115-37.
- Marshall, Dan. (2015) "Humean Laws and Explanation." *Philosophical Studies*. 172: 3145-65.
- Maudlin, Tim. (2007) *The Metaphysics within Physics*. Oxford: Oxford University Press.
- Maxwell, Grover. (1962) "The Ontological Status of Theoretical Entities." In Feigl, Herbert and Maxwell, Grover (eds.). *Scientific Explanation, Space, and Time: Minnesota Studies in the Philosophy of Science*. Minneapolis: University of Minnesota Press.
- Mermin, N. David. (1990) *Boojums All the Way through: Communicating Science in a Prosaic Age*. New York: Cambridge University Press.
- Miller, Elizabeth. (2014) "Quantum Entanglement, Bohmian Mechanics, and Humean Supervenience." *Australasian Journal of Philosophy*. 92: 567-83.
- Miller, Elizabeth. (2015). "Humean Scientific Explanation." *Philosophical Studies*. 172: 1-22.
- Paul, L. A. and Ned Hall. (2013) *Causation: A User's Guide*. Oxford: Oxford University Press.
- Putnam, Hilary. (1975) *Mathematics, Matter, and Method*. Cambridge: Cambridge University Press.
- Redhead, Michael. (1987) *Incompleteness, Nonlocality, and Realism: A Prolegomenon to the Philosophy of Quantum Mechanics*. Oxford: Oxford University Press.
- Roberts, John. (2008) *The Law-Governed Universe*. New York: Oxford University Press.
- Roberts, John. (ms) "Humean Laws and the Power to Explain."
<https://philosophy.unc.edu/files/2013/10/Roberts-DoHumeanLawsExplain-talkversion.pdf>
- Russell, Bertrand. (1914) "On Scientific Method in Philosophy." Reprinted in Russell, Bertrand. (1917) *Mysticism and Logic*. London: Longmans.
- Ryle, Gilbert. (2000) *The Concept of Mind*. Chicago: University of Chicago Press.

- Ryle, Gilbert. (2015) *Dilemmas: The Tarner Lectures 1953*. Cambridge: Cambridge University Press.
- Schaffer, Jonathan. (2003) "Is There a Fundamental Level?" *Noûs*. 37: 498-517.
- Shumener, Erica. (forthcoming) "Laws of Nature, Explanation, and Semantic Circularity." *The British Journal for the Philosophy of Science*. <https://doi.org/10.1093/bjps/axx020>
- Smart, J. J. C. (1968) *Between Science and Philosophy: An Introduction to the Philosophy of Science*. New York: Random House.
- Smart, J. J. C. (1984) "Ockham's Razor." In *Principles of Philosophical Reasoning*, ed. Fetzer, James. Lanham. Maryland: Rowman and Littlefield.
- Sober, Elliott. (1994) *From a Biological Point of View*. Cambridge: Cambridge University Press.
- van Fraassen, Bas. (1980) *The Scientific Image*. New York: Oxford University Press.
- van Fraassen, Bas. (1989) *Laws and Symmetry*. New York: Oxford University Press.
- Wigner, Eugene. (1970) *Symmetries and Reflections*. Cambridge, Massachusetts: MIT Press.
- Wittgenstein, Ludwig. (1965) *The Blue and Brown Books*. New York: Harper.
- Woodward, James. (2014) "Simplicity in the Best Systems Account of Laws of Nature." *British Journal for the Philosophy of Science*. 65: 91-123.