

Toward a Best Predictive System Account of Laws of Nature

Chris Dorst

Abstract

This paper argues for a revised Best System Account (BSA) of laws of nature. David Lewis's original BSA has two main elements. On the one hand, there is the Humean base, which is the totality of particular matters of fact that obtain in the history of the universe. On the other hand, there is what I call the 'Nomic Formula', which is a particular operation that gets applied to the Humean base in order to output the laws of nature. My revised account focuses on this latter element of the view. Lewis conceives of the Nomic Formula as a balance of simplicity and strength, but I argue that this is a mistake. Instead, I motivate and develop a different proposal for the standards that figure into the Nomic Formula, and I suggest a rationale for why these should be the correct standards. Specifically, I argue that the Nomic Formula should be conceived as a collection of desiderata designed to generate principles that are predictively useful to creatures like us. The resulting view—which I call the 'Best Predictive System Account' of laws—is thus able to explain why scientists are interested in discovering the laws, and it also gives rise to laws with the sorts of features that we find in actual scientific practice.

1 Introduction

2 The LOPP

3 A Problem with Lewis's Formula

4 A Pragmatic Account of the Nomic Formula

- 4.1 Informative dynamics
- 4.2 Wide applicability
- 4.3 Spatial locality
- 4.4 Temporal locality
- 4.5 Spatial, temporal, and rotational symmetries
- 4.6 Predictively useful properties
- 4.7 Simplicity

5 Conclusion

1 Introduction

Broadly speaking, David Lewis’s Best System Account (BSA) of laws of nature¹ consists of two main components. On the one hand, there is the Humean base: the totality of particular matters of fact about the universe. And on the other hand, there is a certain formula that gets applied to the Humean base in order to generate the laws. This formula, according to Lewis, involves a balance of simplicity and strength. The laws of nature are, roughly, the regularities of the simplest and strongest systematization of the Humean base.²

Lewis says that this balance of simplicity and strength is derived from actual scientific practice: it is what scientists themselves do when they are trying to figure out the laws of nature. In other words, it is the collection of epistemic principles used by scientists to discover the laws.³ Of course, according to the BSA these are not merely epistemic principles; rather they are constitutive of what it is to be a law of nature. In what follows, I will call the collection of these principles—which scientists use to investigate the laws, and Lewis regards as constitutive of the laws—the ‘Nomic Formula’.

There is a lot to like about the BSA. For example, it renders the laws epistemically accessible, and it does so while incorporating the sorts of criteria that scientists themselves use when they investigate the laws. Better yet, it does all this without appealing to various sorts of ‘anti-Humean whatnots’: there are no necessary connections in the Lewisian ontology. However, there is also a lot to worry about. One source of worry concerns Lewis’s characterization of the Humean base. It has been suggested, for example, that Lewis’s characterization is undermined by our current understanding of quantum mechanics.⁴ It has also been argued that Lewis’s appeal to ‘naturalness’ in his characterization renders it epistemically problematic.⁵ Both of these are serious worries that are worth consideration, but I will set them aside here. My focus instead is on the other main component of his view: the Nomic Formula. My purpose in this paper is to propose a modified Nomic Formula, and with it, a modified Best System Account of laws of nature. Lewis, I think, got it wrong. The Nomic Formula

¹See (Lewis [1973], [1986], [1994]).

²For the sake of simplicity, I will ignore Lewis’s discussion of ‘fit’ in this paper. I will also ignore questions about what happens if two different systems come out tied for the best.

³The fact that the balance of simplicity and strength is meant to be derived from scientists’ actual epistemic practices is not a particularly well-advertised part of Lewis’s BSA. However, he makes it explicit in at least two places:

- 1) ‘I take a [best system] 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. 41).
- 2) ‘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).

⁴For discussion of the problems posed by quantum mechanics, see (Maudlin [2007]) and (Loewer [1996]), as well as some illuminating recent responses by Miller ([2013]) and Boghal and Perry ([2017]).

⁵For example, see (Cohen and Callender [2009]; Loewer [2007]; Demarest [2017]).

does not involve—or at least does not primarily involve—a balance of simplicity and strength.

This paper proceeds as follows. In Section 2, I discuss a useful way of thinking about the Nomic Formula in terms of Ned Hall’s concept of a ‘limited oracular perfect physicist’, or ‘LOPP’. In Section 3, I introduce a novel problem for Lewis’s explication of the Nomic Formula (that is, the balance of simplicity and strength). On the basis of this problem, in Section 4 I develop an alternative explication of the Nomic Formula that (i) avoids the problem of Section 3, and (ii) incorporates desiderata aimed to maximize the laws’ predictive utility. Thus I call the resulting view a ‘Best Predictive System’ account of laws. I conclude in Section 5 by summarizing the advantages of this theory and sketching the work that still needs to be done.

2 The LOPP

Ned Hall ([unpublished]) 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. (Hall [unpublished], 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—indeed, I will introduce a significant one in the next section—even if the unofficial idea is perfectly correct. For the official idea, as it were, sticks it 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 [particular matters of fact] [...] The second guiding idea, then, is roughly that the laws are whatever she says they are. (Hall [unpublished], 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.

There is much more that could be said about the motivations for this LOPP-style Humeanism.⁶ For now, it will be sufficient to note one of the main benefits of this sort of view. By regarding the LOPP's claims as the lawmakers, we are able to make the most possible sense out of our 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. Thus, LOPP-style Humeanism can be seen as starting from a prior commitment to the rationality of our scientific investigations into the laws. It is against this background that my criticisms of Lewis will find a foothold.

⁶Versions of LOPP-style Humeanism are developed in (Miller [2013]) and (Bhogal and Perry [2017]).

3 A Problem with Lewis's Formula

Let us turn now to critically examine Lewis's explication of the standards that comprise the Nomic Formula. As we have already noted, 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.

This characterization of the Nomic Formula is defective for several reasons. For example, Roberts ([2008], pp. 6–10) argues that the practice of balancing simplicity and strength cannot be underwritten by actual scientific practice. For similar reasons, Woodward ([2014]) and Hall ([unpublished]) argue that Lewis's explication of the Nomic Formula will deem certain statements to be laws that intuitively are not. I agree with these objections, but I will leave them aside here, because we have bigger fish to fry.

In my view, the objections posed by Roberts, Woodward, and Hall 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 with the following imaginary conversation:

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... (Lange [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. Imagine that you are navigating a complex 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 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 type (4), for that is by far the most useful to you. By contrast, type (1) doesn't tell you what to do at any point in particular, and types (2) and (3) rely on information to which you, as a person inside this maze, do not have access.

Note, however, that from a God's-eye perspective there is nothing wrong with types (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 types (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 type (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) is navigationally quite useless. Generally speaking, then, 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 with certain pragmatic features that will help them navigate the universe. Thus, their epistemic standards are going to reflect this fact: 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.⁷

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 desiderata into the Nomic Formula. To a nonreductionist, 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. To illustrate why, I am again going to draw on some remarks by Ned Hall ([unpublished], pp. 39–40). Hall suggests that, given a reductionist metaphysics, it actually makes a great deal of sense to think of the laws pragmatically.

Hall's central point is that, given the reductionist's fundamental metaphysical picture of the world—be it a collection of particles moving around, or space-time points instantiating perfectly natural properties, or the more vaguely characterized 'particular matters of fact' that I have been using—there is a challenge that does not show up for nonreductionist 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' ([unpublished], 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 explana-

⁷It 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. 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.

tion of the laws' attention-worthiness, something that doesn't appeal to their metaphysical status.

As Hall suggests, the most plausible explanation of 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 in the first place. So the appeal to pragmatism is in fact quite appropriate for the reductionist, and it suggests the line of thought that I am going to pursue below. That is, I am going to develop the idea that we seek knowledge of the laws because it provides practical advantages.

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.⁸ 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.⁹ Alternatively, it's possible that the information we care about won't get swamped out, but instead it will

⁸This is not to say that there is no use in studying that sort of information, only that that information is not particularly useful when considered as an end in itself. For example, cosmologists focus on information about events that occurred long before we were all born, but surely they would bristle at the suggestion that their work is useless to us. Rather, they would probably suggest that by studying the conditions of the early universe, we are able to test and modify the principles that we use around here. For if those principles didn't apply in the early universe, then this suggests that there could be something subtly incorrect about them, such that they may lead us to make incorrect predictions even when we apply them to conditions around here. In other words, there can be practical uses for studying the subject matter of cosmology, even if the practical uses of the content of that subject matter are themselves minimal. (Thanks to an anonymous referee for pressing this point.)

⁹Case 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.

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 that would be of very little utility to you, such as telling you the total number of lefts and rights that you have to take. 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—that is, they might give us sketches where we need details, or they might give us details, but in a form that is utterly useless 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 quite useless to 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 this: 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. For example, the laws support counterfactuals, figure in scientific explanations, underwrite inductive inferences, and so on. Of these 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 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', and to say that 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 system with 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 . Then it would obviously be quite useful to have 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, that is, if it is deterministic. If instead it only gives us probabilistic statements about 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 I will discuss some reasons it could fall short of determinism in Sections 4.3 and 4.6.

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, as a referee pointed out, 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 here 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

‘phase space’), and (2) a dynamical principle (or set of principles) that specifies the trajectory of any given system through phase space.¹⁰ 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.¹¹

4.2 Wide applicability

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 types of systems we want to make predictions about. In other words, we should 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 among these systems. Specifically, 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. Consequently, our 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.¹²

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 the-

¹⁰In classical mechanics, the phase space of an n -particle system is a $6n$ -dimensional space, and the dynamical principles are deterministic. In quantum mechanics, the phase space is a vector space (whose dimensionality varies, but is often infinite), and the dynamical principle is the Schrödinger equation, which again is deterministic.

¹¹Here I am in agreement with Hall ([unpublished], p. 1) and Maudlin ([2007], p. 13), who both regard the dynamical implications of the laws as a defining characteristic.

¹²Here 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.

orists, such as Hall ([unpublished]), 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.¹³ 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 (that is, 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 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.¹⁴

¹³In 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.

¹⁴Of 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

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 (for example, 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, 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 nonlocality inherent in quantum mechanics was demonstrated by Bell’s theorem, along with the subsequent experimental violations of Bell inequalities.¹⁵ 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

laws fail to make reference to n , but that none of them make any reference to it at all.

¹⁵See (Aspect *et al.* [1981]).

locality desideratum (even the revised one), and thus you might think that it spells trouble for my account of the Nomic Formula.¹⁶

However, there are three points to note. First, and most basically, remember that the extent to which we can design a theory that satisfies these various desiderata is dependent on how kind nature is. And in fact, I think we really ought to expect not to be able to satisfy all of the desiderata simultaneously, because some of them pull in opposite directions. Consider just the two desiderata of an informative dynamics and spatial locality. The latter desideratum directs us to use only local variables to 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, it is still telling that our earlier theories (such as 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 failure of quantum mechanics 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 fact that a radically nonlocal theory would be predictively useless to creatures like us, and by 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. That is, it does not leave us in a situation of total predictive ignorance, because it is not radically nonlocal. We are still able to use quantum mechanics to make reliable predictions even without knowing what is going on at very distant locations. This is because entanglement phenomena, which are the source of nonlocality within the theory, are exceedingly fragile, so oftentimes they have little to no effect on the behavior of systems we care about.¹⁷

¹⁶Thanks to an anonymous referee for making me appreciate the significance of this worry.

¹⁷The precise reason that entanglement is so fragile is controversial, and different interpretations of quantum mechanics explain it differently. Relatedly, an anonymous referee has suggested that the very existence of competing interpretations of quantum mechanics may be explicable, in an interesting way, on the view I am developing here. Roughly, the thought is that quantum phenomena are so strange that they render it increasingly likely that conflicts will arise between these various desiderata. When we cannot satisfy all of the desiderata at once, we are forced to choose which ones to weight more heavily. 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, though they arguably do so at the cost of others (such as those mentioned in Sections 4.6 and 4.7). Obviously, this is an issue that deserves more exploration than I can give it here, but I think the basic idea is highly suggestive.

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 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 (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 (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.

Note, also, 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.¹⁸

¹⁸Similar remarks would apply to many of the other desiderata that I am discussing here: Lewis's theory has no easy way of explaining why the actual laws satisfy them to the extent that they do. As I mentioned in Section 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.

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). 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. Specifically, 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.¹⁹

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, 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 fancy footwork. 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.

¹⁹For 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.

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 infer that, four minutes ago, the atom had 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 (Lewis [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, ‘ $\forall xFx$.’ 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 value simplicity in our scientific theories. Consider, for example, how it 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, claiming that ‘You can recognize the truth by its beauty and simplicity’ ([1965], p. 171). He of course grants that a theory has to have the right sorts of observable consequences to be deemed correct, but he thinks 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 can tell that they are not true, he says, because of their complexity. 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 dynamical implications
- 2) Applicability to a wide range of different systems
- 3) Spatial locality (that is, most systems of interest are treated as enclosed or quasi-enclosed)
- 4) Markovian dynamics that proceed from past to future
- 5) Spatial, temporal, and rotational symmetries
- 6) Properties and magnitudes that are predictively useful (that is, they are ascertainable at a time when they can still be used to make predictions)
- 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 precisely this collection of features should have coalesced under the concept of a law of nature.

5 Conclusion

By way of concluding, I want to provide a sketch of the overall theory I am advancing and point out the work that still needs to be done.

The overall architecture of my theory is the same as the original BSA: the laws of nature are the results of applying the Nomic Formula to the Humean base. I have not taken up the question here of exactly how to conceive of the

Humean base. Rather, I have focused on critiquing Lewis's conception of the Nomic Formula, and on developing a better one. In accordance with LOPP-style Humeanism, the Nomic Formula is conceived as the epistemic standards physicists use to figure out the laws of nature. Furthermore, according to my theory, these standards consist of a variety of desiderata designed to output principles that are 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 physicists' epistemic practices. Essentially, it elevates the epistemic standards used by physicists to standards that are constitutive of lawhood, thereby obviating the need to provide an independent epistemic justification for those standards. Third, and most significantly, it makes sense of why physicists 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 both epistemic and pragmatic grounds. That is, 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.

Of course, the work is not yet done. As I mentioned before, there are outstanding problems with Lewis's characterization of the Humean base: it looks to be inconsistent with quantum mechanics, and its appeals to naturalness make it epistemically problematic. Without an adequate account of the Humean base, we are short roughly one half of a theory; we have our revised Nomic Formula, but we do not have the raw material to which it is supposed to be applied. It is well beyond the scope of this paper to take up this issue here. But as I see it, in order to complete a reductive theory of laws along the lines suggested in this paper—a 'Best Predictive System Account', if you will—one would have to conjoin the above account of the Nomic Formula with a revised conception of the Humean base that avoids the problems that plagued Lewis's. For now, I will say only that Cohen and Callender ([2009]) appear to make some promising headway in that regard. I expect that any tenable account of the Humean base will have to proceed using their discussion as a basis.

To wrap things up, I want to consider a highly suggestive quotation from *Fact, Fiction, and Forecast*. There, Nelson Goodman wrote of the 'Humean idea that rather than a sentence being used for prediction because it is a law, it is called a law because it is used for prediction' ([1983], p. 21). Granted, Goodman viewed this only as a 'first approximation', not a complete theory of lawhood. The natural way to pursue this suggestion would be to think carefully about what sorts of principles (or 'sentences') creatures like us would find predictively useful, and then determine how closely those principles resemble actual putative laws of nature. The more closely those principles resemble the laws, the more plausible Goodman's suggestion becomes. If my arguments above are correct, then, they constitute a vindication of Goodman's idea.

Acknowledgements

I would like to thank Marc Lange, Ned Hall, Matt Kotzen, John Roberts, Kevin and Stanley Dorst, and two anonymous referees for their many helpful comments and suggestions on earlier drafts.

*University of North Carolina, Department of Philosophy
Chapel Hill, USA
dorst@live.unc.edu*

References

- Albert, D. and Loewer, B. [1988]: ‘Interpreting the Many Worlds Interpretation’, *Synthese*, **77**, pp. 195–213.
- Aspect, A., Grainger, P. and Roger, G. [1981]: ‘Experimental Tests of Realistic Local Theories via Bell’s Theorem’, *Physical Review Letters*, **47**, pp. 460–3.
- Bhogal, H. and Perry, Z. [2017]: ‘What the Humean Should Say About Entanglement’, *Noûs*, **51**, pp. 74–94.
- Cohen, J. and Callendar, C. [2009]: ‘A Better Best System Account of Lawhood’, *Philosophical Studies*, **145**, pp. 1–34.
- Demarest, H. [2017]: ‘Powerful Properties, Powerless Laws’, in J. Jacobs (ed.), *Putting Powers to Work: Causal Powers in Contemporary Metaphysics*, Oxford: Oxford University Press, pp. 38–56.
- Einstein, A. [1948]: ‘Quanten-Mechanik Und Wirklichkeit’, *Dialectica*, **2**, pp. 320–4.
- Feynman, R. [1965]: *The Character of Physical Law*, New York: Modern Library.
- Gell-Mann, M. [1964]: ‘A Schematic Model of Baryons and Mesons’, *Physics Letters*, **8**, pp. 214–5.
- Goodman, N. [1983]: *Fact, Fiction, and Forecast (Fourth Edition)*, Cambridge, MA: Harvard University Press.
- Hall, E. [unpublished]: ‘Humean Reductionism About Laws of Nature’.
- Lange, M. [2009]: *Laws and Lawmakers*, New York: Oxford University Press.
- Lewis, D. [1973]: *Counterfactuals*, Malden: Blackwell.
- Lewis, D. [1983]: ‘New Work for a Theory of Universals’, *Australasian Journal of Philosophy*, **61**, pp. 343–77.
- Lewis, D. [1986]: *Philosophical Papers, Volume II*, New York: Oxford University Press.
- Lewis, D. [1994]: ‘Humean Supervenience Debugged’, *Mind*, **103**, pp. 473–90.
- Loewer, B. [1996]: ‘Humean Supervenience’, *Philosophical Topics*, **24**, pp. 101–27.
- Loewer, B. [2007]: ‘Laws and Natural Properties’, *Philosophical Topics*, **35**, pp. 313–28.

- Maudlin, T. [2007]: *The Metaphysics within Physics*, New York: Oxford University Press.
- Miller, E. [2013]: ‘Quantum Entanglement, Bohmian Mechanics, and Humean Supervenience’, *Australasian Journal of Philosophy*, **92**, pp. 567–83.
- Roberts, J. [2008]: *The Law-Governed Universe*, New York: Oxford University Press.
- Wigner, E. [1970]: *Symmetries and Reflections*. Cambridge, MA: MIT Press.
- Woodward, J. [2014]: ‘Simplicity in the Best Systems Account of Laws of Nature’, *British Journal for the Philosophy of Science*, **65**, pp. 91–123.