HUMEANISM AFTER THE BSA: A SECOND BEST SYSTEM

by

MICHAEL TOWNSEN HICKS

A dissertation submitted to the

Graduate School-New Brunswick

Rutgers, the State University of New Jersey

In partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

Graduate Program in Philosophy

Written under the direction of

Barry Loewer

And approved by

____

____

____

____

____

New Brunswick, New Jersey

JANUARY 2016

ABSTRACT OF THE DISSERTATION

Humeanism After the BSA: a Second Best System

By MICHAEL TOWNSEN HICKS

Dissertation Director:

Barry Loewer

This dissertation explores the idea that laws of nature are tools for gaining and employing information about the world.

In my first chapter, I focus on the Best System Account of laws of nature, according to which the laws are those generalizations which form a deductive system best combining strength and simplicity. I claim that we must focus on the applicability of laws to the sorts of subsystems of the universe with which we interact and manipulate. This provides us with the tools to develop new notions of informativeness and simplicity, and sheds light on our preference for dynamic explanations.

In the second chapter, I show how best to extend a Best Systems based account to the special sciences. I argue that extant theories of the relationship between fundamental and special sciences either deemphasize interscientific connections or do little justice to the counterfactual robustness of laws in the special sciences. We need the special sciences in

addition to physics because often we have only coarse-grained information about our surroundings, so we need coarse-grained laws to make predictions. I provide a precise characterization of coarse-graining and show how it can be applied to the relationship between the special sciences and physics.

In the third chapter, I argue that importing notions from epistemic utility theory provides an account of objective chance which explains both the connection between chances and frequencies and the connection between chances and degrees of belief. This account relies crucially on the notion of an isolated subsystem developed in the first chapter of the dissertation, and aims to show that the objective chances we find in science are fit to act as a guide to our beliefs. I show how measures of accuracy currently being developed in epistemic utility theory can be applied to measure how well the chances of some class of events–say, coin tosses–fit the frequency of outcomes.

DEDICATION

This dissertation is dedicated to my parents, Donald Walter and Celeste Harrison Hicks, who encouraged me to write it and who halfway lived to see it completed.

Thanks to my advisor, Barry Loewer; my committee, Ned Hall, Branden Fitelson, Brian McLaughlin, and Jonathan Schaffer; and everyone who gave me comments, suggestions, and feedback, including Erik Hoversten, Thomas Blanchard, Heather Demarest, Zee Perry, Harjit Bhogal, Marco Dees, Eddy Keming Chen, David Albert, and Alison Fernandez.

Contents

Introduction

The world as we experience it is shot through with possibility and necessity; facts about what can happen, what must happen, and what would happen (if something else did) are as much a part of our picture of the world as lampshades, rocks, or the color green. Judgments about how things could go or would have gone are made as naturally as judgments about what has happened or will happen. But facts about what would have happened are not like facts about what did happen: they discribe ways the world could have gone but didn't. There is only one world; it only goes one way. Why do we care about ways the world could have been?

Here, I seek to answer this question in a way compatible with two constraint: Humeanism and actualism. Humeanism is the view, named for David Hume, that there are no necessary connections between distinct existences: that all there is to the world are particular things, their properties, and their relations to one another. Actualism is the view that there is nothing beyond the actual world. Taken together, these imply that *modal* facts about what must happen, what could happen, and what would happen are somehow be grounded in or made true by *categorical* facts: facts about what does happen.

These constraints are adopted not just because, by adopting them, we arrive at a more austere and aesthetically pleasing metaphysics (although of course we do). Rather, they are adopted because explanations are best when they explain the confusing in terms of the familiar. The best hope for explaining the importance and utility of *modal* notions–notions about what could, must, and would happen–explains these things in terms of what does happen. Our metaphysical account of modality ought to lend itself to an account of the importance of modal notions in our practices of prediction, explanation, and decision making. This cannot be done by stipulation; positing *sui generis* properties or relations of necessitation and then asserting that they have important practical and epistemic import is less than unsatisfying–it is an abdication of our responsibility as metaphysicians. We must connect these notions to those things that we are interested in, and explain how our interests give rise to them. I believe that humeanism provides us with the best framework in which do offer such an explanation.

This dissertation approaches the mystery of modality through the laws of nature. Laws of nature constrain what can happen: only those things compatible with the laws are possible, and if something is implied by the laws, it must happen. The laws also determine the counterfactual structure of the world; when we try to figure out whether, if I had gotten on the plane, I would be there by now, we imagine a world different only in that I got on the plane, and use the laws to determine whether I would be there by now. This connection between laws and counterfactuals is most apparent in the use of controlled experiments to discover the laws. When constructing experiments, we are interested in what *would* happen if the variables we can control *were* different; we investigate this by making them different, and seeing what does happen. We then check this against what was predicted by our laws. If the laws are robust enough to hold in a variety of actually different situations, we infer that they would hold if things were not as they actually are.

Here, I explore the idea that laws of nature are tools for gaining and employing information about the world. The view defended here is a development of the regularity view of natural laws, which has its roots in Hume, and whose prominent proponents have included John

Stuart Mill, Frank Ramsey, and David Lewis–the poster adult of modern Humeanism. This view of laws is called 'the regularity theory'.

The regularity theory of natural law holds that laws of nature are merely regularities; they are just things that happen. They are special because they happen in a pattern. They are predictable and they are uniform, but ultimately, they are merely happenings. The central challenge for this view, then, is to explain their air of necessity. If the laws of nature merely express the way things are, why does it seem that things *must* comport with them? Why do we say that things *have to* obey the laws? And what separates the laws from merely accidental regularities–the things that could have gone either way?

Answering this question is important not merely because the laws seem to us to be necessary; if that were the case, then an acceptable response would be not to explain their necessity but instead to explain its appearance away. But the necessity of the laws justifies their central place in scientific inquiry. The laws justify our predictions of future events; they tell us what to expect to see. We can expect the world to obey the laws because it must do so. Our explanation of the 'must' needs to go some distance towards explaining why we should expect the laws, but not the accidental generalizations, to hold in the future.

Understanding the necessity of the laws is also important because doing so can help us understand counterfactuals: the laws are what make it the case that if I *were* to leap from my window I *would not* fly, that if I *were* ten thousand leagues under the sea I *would* be crushed by the pressure, and that if I *were* to ride my bike in front of an oil truck I *would* suffer serous injury. These counterfactuals, which are made true by the laws, are central to my decisions *not* to leap from the window, deep sea dive, or bike in front of an oil truck, despite the fact that their antecedents and consequents (hopefully) are never actually realized. So we must explain the necessity of the laws in way that sheds light on their connections to counterfactuals and decision theory.

The regularity theory's answer to this question is pragmatic: the laws are necessary because it is useful for us to hold them fixed when making inferences about what has happened, what will happen, and what would happen. The most prominent regularity theory, which was first expressed by John Stuart Mill, but which has been modified and defended by Frank Ramsey and David Lewis, takes the laws to be necessary because they are the fewest axioms from which we could deduce everything (or anyway, a lot). The idea is simple: we are deductive creatures. Our memory is limited, but our capacity for inference is not. Hence, if there are some short, relatively simple sentences from which we can infer many truths, our best bet is to take those and infer further truths from them as needed. When making predictions or evaluating counterfactuals, we hold the laws fixed so that we can use them in deducing what will or would happen. Their necessity comes from this feature: they are the facts we *should* hold fixed when considering what will or would happen.

This view of laws is called the Best Systems Account (BSA): the laws are the axioms of that deductive system which best balances strength (inferential power) with simplicity (shortness, or tractability for creatures like us). I discuss this view of laws at length in ch. [lawsandroles], §[BSA] This dissertation seeks to improve upon and extend the BSA in three ways, addressed in three separate chapters.

First, the BSA has a notorious problem explaining the fact that the laws are typically dynamic. Most of the laws we find in the sciences tell us how systems (of particles, ecosystems, and economies) go from one state to another. Rarely do they provide constraints on what sort of systems there are, or (even worse) just tell us about the global state of a system at some time. But the BSA tells us that a statement earns the status of law by being short and highly informative. This leaves room for laws which tell us not how things develop but how things are. This is not a bullet we can bite: the sciences are focused on figuring out how systems develop. We seek to explain, so far as we can, *why* the sciences work the way they do. If we cannot explain why they look for that sort of information, we have failed.

The first chapter of this dissertation (ch. [lawsandroles]) aims directly at this problem, and provides a regularity theory which correctly differentiates laws from mere accidents. This theory, the Epistemic Role Account of natural law, takes laws to be those generalizations best fit to sit between induction and deduction. They must be learned by limited agents in the world, and so cannot be facts only about the world as a whole; they must be useful to those agents, and so must provide us with information we can deploy as needed.

I use this sketch of the role of laws to identify virtues a law-system must have: it must be broad, in the sense that it can be used to model a variety of subsystems of the world. It must be informative. Its parts must be independently testable. I argue that lawsystems that play this role and have these virtues will consist primarily of dynamic information. Thus, I argue, the proponent of the regularity theory can account for the fact that laws are typically dynamic without sacrificing her metaphysical scruples.

Ch. [lawsandroles] focuses on deterministic, fundamental laws. The remainder of the dissertation seeks to extend the view of laws outlined in the first chapter to nonfundamental law systems (ch. [Coordination]) and probabilistic lawsystems (ch. [fitfit]).

Most theories of laws, and of the necessity of laws, focus on fundamental laws: laws which govern or describe the actions of the most fundamental properties of the world. Presumably, because everything is determined in some way or another by the actions of the fundamental stuff, be it particles or fields or strings, the doings of biology, chemistry, and economics are determined by the doings of physics. So, one might think, if we have an explanation of the necessity of the fundamental, we automatically have a theory of the necessity of the higher-level sciences. Or, one might worry, if the workings of physics are determined by the laws of physics, and the workings of biology are determined by the workings of physics, then the workings of biology are too determined by the laws of physics: biological laws are mere epiphenomena.

In ch. [Coordination] I argue that one would be wrong if one one thought or worried either of these things. The necessity of the laws of biology is independent of the necessity of physics. Each of the sciences studies epistemologically independent phenomena; each science independently discovers pragmatically useful generalizations. I argue that previous accounts of the relationship between the sciences either overemphasize the importance of physics and thereby fail to account for the independence of the special sciences, or overestimate their independence, and thereby fail to account for the fact that the phenomena the special sciences study are grounded in the phenomena studied by physics. I argue that the regularity theory I defend is ideally suited to balance this metaphysical dependence with epistemological and methodological independence.

The lawsystems I discuss in chs. [lawsandroles] and [Coordination] are deterministic. But not all worlds are deterministic worlds; some are instead governed by objective probabilities. Just as I hold that laws are encode pragmatically interesting information, rather than describing some *sui generis* entities or fundamental necessity, I hold that probabilities are primarily useful to us because of the information they bear about the world. In ch. [fitfit], I show how an understanding of objective probabilities fits into the view of laws described in this book. I argue, following Lewis, that these are best understood via the Principal Principle, which links objective chance with our subjective degrees of belief. I then argue that objective chances should be understood as the most accurate degrees of belief possible for agents subject to the same sorts of constraints we are: namely, those who are relegated to a limited place in space and time, who must make judgments based on the local, qualitative features of their surroundings.

These chapters, taken together, represent a serious modification of the orthodox humean view of laws. They challenge not just the details of traditional humeanism, but also its motivating story about the importance and role of laws. They hang together by showing how a new conception of the role and importance of laws of nature can answer outstanding questions about the relationship between laws and initial conditions, the relationship between laws in different scientific disciplines, and the relationship between objective probability, the world, and our beliefs. If, as I argue, the humean can give a good account of each of these features of lawhood, she will have a good start in explaining the central place of modality in our experience of the world.

Ch. 1: Laws and Epistemic Roles

Introduction

What are laws of nature? They are universal generalizations. But not all universal generalizations are laws. Laws are those generalizations which underwrite counterfactuals, license predictions, feature in explanation, and are primarily discovered only via empirical inquiry[1]. We can appeal to examples: Schrödinger's equation, Einstein's field equations, Newton's equations of motion. And we can give a philosophical account in terms of necessary and sufficient conditions. In this chapter, I will examine one of the most prominent philosophical theories of natural law: the Mill-Ramsey-Lewis Best System Account (hereafter 'the BSA'). I will show that, in its current iteration, it fails to correctly differentiate between laws and non-laws. I'll open by showing the commitments of the view, in its now canonical form: Lewis' Best System Account of law. I'll then present two arguments against it.

I will then clearly diagnose why the BSA falls short: the BSA puts too little focus on induction. Laws must be discovered empirically by limited agents operating locally. Recognizing this allows us to more accurately explicate the theoretical virtues scientists seek in laws.

I will not address criticisms from proponents of more metaphysically robust accounts of laws, which attempt to show that the mere generalizations cannot play some of the that

[1]This separates natural laws from the laws of mathematics.

laws are supposed to play, such as supporting counterfactuals, underlying explanations, and supporting predictions. Instead, I'll be arguing that the current incarnation of the regularity theory cannot cleave the laws from the accidental truths. However, some external critics have argued that the BSA's reliance on strength and simplicity don't connect with the practice of science so the regularity theory must be abandoned. I hope here to show that a modified regularity theory can overcome this difficulty.

In this chapter I focus on the difference between dynamic and static generalizations in fundamental laws; I present a regularity theory which can account for this. In ch. 2 and ch. 3, I will appeal to the account of the role of laws that motivates this view in developing it to overcome other obstacles for the regularity theory.

The dissertation has implications which go far beyond the regularity theory, for the virtues I discuss generalize. Typical discussions of these virtues, at least amongst metaphysicians of laws of nature, focus on simplicity and strength even in metaphysically robust regimes. It is this paradigm that I attack in this paper. Here I seek to extend the short list of features that scientists should pay attention to. This extension, I believe, is warranted even for those who take these virtues to be epistemic guides to the the laws rather than metaphysical determiners of them.

The Best System

The regularity theory holds that laws of nature are merely generalizations. The 'mere' here is doing a hefty bit of work: it distinguishes the laws of the regularity theorist from the more metaphysically robust laws of her interlocutors, who I will call 'modalists'[2]. Modalist views are less metaphysically perspicuous than the regularity theory because they claim that, to be laws, a generalization must be backed, made true, or associated with a relation between properties (Armstrong (1983), Dretske(1977), Tooley (1977)), the essences of properties (Shoemaker (1980), Ellis (2001), Bird (2007)), *sui generis* facts about production (Maudlin (2007)), or irreducible counterfacts (Lange (2007)). These bits of metaphysical machinery are meant to underwrite the necessity of the laws. Although each of these metaphysical machines is distinct, modalists are united in holding that facts about nomic necessity, or modal facts of some sort, are fundamental.

In contrast to these views, the regularity theorist holds that the laws are not backed or made true by anything beyond their instances and they are made *laws* by nothing more than the sum of non-nomic facts at a world. What makes them special is not some metaphysical fact, but instead our epistemic interests. A law statement, like $F = ma$ is made *true* by its instances; it's made a *law* by the total distribution of fundamental properties.

By abandoning any attempt to metaphysically explicate what David Hume called 'necessary connexions', the regularity theorist must find something to do the work these connections do for the modalist: cleave the laws from the accidentally true generalizations. This must be done without introducing any primitive modal machinery, else the regularity theorist will find herself a converted modalist.

[2] I'm using this term, rather than 'anti-Humean', partially because a regularity theorist need not *reject* irreducible modality–she is merely not committed to it by her theory of laws. The regularity theory may be true though Humean Supervenience is false, as is argued by Demarest (forthcoming).

But this isn't all she must do. For a philosophical account of laws ought to be explanatory. It should not merely tell us which generalizations are laws; it should tell us *why* generalizations fitting its profile are fit to play the role of laws. Laws support counterfactuals, they underwrite predictions, and they are suitable bases for induction. The regularity theorist's account of *what* separates the laws from the non-laws ought permit us to tell a believable story about *why* we take generalizations with those features to be special. If she does so, she will have a leg up on the modalist, who must tack these *epistemic* features onto her *metaphysical* posits.

Orthodox Humeanism

The regularity theorists' answer to this challenge has been in circulation since John Stewart Mill's *A System of Logic*. Here's Mill:

> According to one mode of expression, the question 'What are the laws of nature?' may be stated thus: 'What are the fewest and simplest assumptions, which being granted, the whole existing order of nature would result?'

Mill tells us that the laws are the sparcest set of truths from which we can derive everything. Similar characterizations of law can be found in Ramsey (1928)[3] and, in its current form, in David Lewis (1983).

Lewis characterizes laws as the general axioms of whatever deductive system best combines simplicity, informativeness, and probabilistic fit. The BSA has changed slightly since Mill's 1843 explication: modern regularity theorists have backed off of Mill's claim

[3]Like Lewis, Ramsey slightly modified Mill's view (for example, Ramsey counts statements derived from laws together with robust initial conditions as laws, but reserves the term 'fundamental laws' for those statements that feature directly in the best axiomatization of facts). Ramsey later rejected this view of laws (Ramsey (1929)).

that the laws need imply *everything*–to do so would unnecessarily make everything nomically necessary. Instead, they take laws to be informative but compatible with a sphere of nomic possibility. While work has been done to nail down these notions of simplicity and informativeness, and to extend the view to probabilistic systems, little more has been added or subtracted from the core of Mill's view.

Contemporary regularity theorists who endorse versions of the BSA include Barry Loewer (Loewer (1996, 2007), Helen Beebee (2000, 2006), Craig Callender and Jonathan Cohen (2009, 2010), and Ned Hall (MS). Here I'll open with Lewis' canonical account and then show how these theorists modify the BSA.

According the BSA, when we're generating and evaluating a system of laws, our goal is to unify and maximize our knowledge. So we need to find a system in which a few statements imply a lot–ideally, everything. This leads us to recognize two virtues of systems, which weigh against one another. First, the system must be *strong*: it must imply a lot. Characterizing this virtue is tough. If two systems both fail to imply everything, there's no non-arbitrary way to measure which one implies more[4]. Lewis claims that the strength of a system should be measured by the number of possible worlds it rules out. But any two systems will equally rule out infinitely many worlds. So we will be able to compare the strength of two systems if and only if (a) one system excludes a subset of the worlds the other excludes, or (b) we have some way of constructing a preferred measure over worlds.

[4]This follows from the fact that there are infinitely many propositions implied by each system. Though measures can be introduced over the set of implications of the laws, there is no unique measure, and a measure can be concocted favoring any lawsystem.

Strength is mirrored by *fit* for systems containing probabilistic laws. Systems fit a world better when they give the world a higher probability[5] Systems with a high fit give accurate probabilistic information: if we match our credences to the chances given by the system, we'll have high credence in truths and low credence in falsehoods. I discuss *fit* in more detail in ch. 3. Finally, the system must be *simple*. Understanding simplicity is a notorious problem for the BSA: specifically, any syntactic account of simplicity will be language dependent. We should worry if a gruesome system is more simple by linguistic fiat. In considering this, Lewis presents the following counterexample:

> Given a system *S*, let *F* be a predicate that holds at all and only the worlds were *S* holds. Take *F* as primitive, and axiomatize *S* (or an equivalent thereof) as ∀xFx. If utter simplicity is so easily obtained, the ideal theory may as well be as strong as possible. [...] Then, after all, every regularity will be a law. This must be wrong.
>
> (Lewis (1983))

The problem presented by the predicate F is straightforward: if our choice of language is maximally free, then the simplicity requirement is toothless. In response to this, Lewis restricts the language in which the the system can be couched: the language must include only predicates which refer to perfectly natural properties. Call this 'the naturalness constraint'. According to Lewis, we discover these natural properties empirically: one of the jobs of physics is to come up with a list of the fundamental properties.

This, then, is the orthodox BSA. The laws of nature are those generalizations in the set of truths which jointly maximizes

[5]Fit is similarly hard to measure, as if a world contains infinitely many events (like ours probably does) it will be assigned probability zero.

- STRENGTH, which measures the deductive informativeness of the laws.
  - System L is *stronger* than system L* if and only if L rules out more worlds than L*.
- FIT, which measures the probabilistic informativeness of the laws.
  - System L has a *higher fit* than system L* if and only if L gives the actual world a higher probability than L* does.
- SIMPLICITY, which measures the simplicity of the laws.
  - System L is *simpler* than L* if and only if the sentences of L, when written in a language whose predicates correspond to perfectly natural properties, are syntactically shorter than those of L*.

It's important to remember that a satisfactory account of lawhood will not merely supply us with necessary and sufficient conditions for lawhood. It will also tell us why we should care about those generalizations which meet those conditions. The proponent of the BSA cannot merely identify those virtues–simplicity and strength–which a lawbook must maximize to qualify as the laws of the world. They must also explain why we should pay attention to systems with those virtues.

Such a story can be gleaned from BSA proponents: according to the BSA, what we seek in laws are efficient organizational tools for our knowledge. So says Ramsey:

> Even if we knew everything, we should still want to systematize our knowledge into a deductive system. [...] As it is, we do not know everything; but what we do know we tend to organize as a deductive system and call its axioms laws.

Ramsey (1928).

Another good illustration is the T-shirt analogy: suppose you encounter God, and She offers to give you any information you desire. You ask for everything. God, obligingly, starts listing facts from the beginning: ``There is a particle with mass m_1 at location $< x_1, y_1, z_1, t_1 >$, and a particle with mass m_2 at point $< x_2, y_2, z_2, t_1 >$ and another particle at...“ You quickly realize God is going to list far too many facts for you to remember when you get home and you have nothing to write on except your t-shirt. So you ask God, ``O Almighty, could you give me a short enough summary to fit on this t-shirt?” ``Ah!“ says God, ``that’s a different question entirely!” And She thinks for a minute, and then gives you a few sentences that sum everything up. These are the laws.

Here we have two explanations of the usefulness of laws. In the second, illustrated by the t-shirt analogy, we seek laws because of our cognitive limitations: we are just too stupid to find a list of all facts useful, but, since we’re pretty good at deduction, a summary consisting mostly of generalizations is both retainable and utilizable to us. The first, described by Ramsey, sees laws as an organizational tool. Given these conceptions of laws, the virtues of simplicity and strength are quite natural: organizations of knowledge are better when they are *more organized*–that is, simpler–and they are better when they *organize more*–that is, they are stronger. Any attempt to modify this short list of virtues had better show that the purported modifications further this role of laws, or successfully argue that Mill, Ramsey, and Lewis have misconstrued the role of laws in our epistemic lives. I set myself to this latter task in §[Laws and Roles]. But first, I’ll look at some current modifications of the traditional BSA.

Heterodox Humeanism

Now that we've seen orthodoxy, let's discuss heterodoxy: there are two current regularity theories which dispense with the naturalness constraint: Barry Loewer's Package Deal Account, and Callender and Cohen's Better Best System Account (Callender and Cohen, (2009, 2010)).

I'll examine the Package Deal Account first. According the the Package Deal Account, each law system is allowed to generate its own language. The simplicity and informativeness of the system is then judged by the ease with which the language can be used by us.

One way to cash this out is as follows: each candidate lawbook comes up with a new set of terms and a translation manual between those terms and an ordinary natural language for discussing middle-sized dry goods. The informativeness of the system is judged by the set of statements in the natural language which can be derived from it. The simplicity of the lawbook is judged by *both* the length of the statements in the lawbook's own language and the length of the translation manual between the lawbook's language and the natural language[6].

The Better Best System account follows a different path in rejecting the naturalness constraint. According to Callender and Cohen, each science comes with its own language,

[6]Although this translation manual scheme will be language dependent, we can assume that our languages have enough similarities that this won't be an issue. This does make the laws more anthrodependent than some might like; however, any natural language, human or alien, would be couched in terms of middle-sized dry goods, and so any language sufficiently different from ours to deliver different laws would belong to beings with a different enough conceptual scheme that we would have independent reason to doubt that they have a concept of law.

which describes its domain of inquiry. Although all languages describe the same world, none does so in a more or less natural way. The categorization scheme of biology is, on this view, just as natural as the categorization scheme of physics. Each systematizes different features of the world, which may be more or less interesting to us for a variety of pragmatic reasons. The predicate F isn't ruled out because it's not a natural kind. Rather, there is a science whose sole law is $\forall x F x$, but this science isn't in any way interesting to us. In addition to dropping the naturalness constraint, some regularity theorists drop or modify other features of the Best System Account. Loewer and David Albert argue that a probability distribution over initial conditions is a law, despite the fact that it is not a generalization[7].

I am sympathetic to all of these alterations. Loewer and Albert point out that restricting laws to generalizations isn't motivated by the BSA's view of laws as organizational tools, and all of these authors seek to make laws more accessible to us–and so better organizations of the facts we care about. In what follows I will be neutral about the naturalness constraint and argue that the generality constraint is not justified. But I think these alterations are not enough. Here I commit a more extreme heresy: the virtues of simplicity, strength, and fit, as understood by Lewis, are not the only virtues which determine which system is best.

[7]Lewis allows that the best system may contain statements that are not generalizations, but holds that they are not laws. Albert and Loewer reject this.

The Best Is Not Good Enough

The BSA holds that being a part of the simplest, strongest, and best fit system couched in perfectly natural terms is necessary and sufficient for being a law of nature. This is false. Being a member of the simplest, strongest systematization at a world is neither necessary nor sufficient for being a law of nature. In §[not good enough] I will present arguments against the BSA. The first (§[boundary conditions]), the argument from boundary conditions, shows that the BSA lacks the resources to distinguish between laws and boundary conditions. The second, (§[scientific practice]) an argument from scientific practice, shows that the virtues articulated by the BSA are not those sought by scientists, and that the BSA's explanation of the role of laws cannot adequately explain the norms on theory choice scientists do respect. In §[Example] I provide a counterexample to the BSA illustrating these failures, and in §[solutions] I conclude by arguing that extant solutions are unmotivated by the BSA's account of the role of laws.

Laws and Boundary Conditions

Some statements are part of the simplest, strongest systematization of a world, but are not laws at that world. Some of these can plausibly be construed as law-like boundary conditions; others seem to be merely contingent truths. The first we will consider is due to Ned Hall:

> Suppose, for example, that there is some moment of time such that [...] there is some relatively simple, compact way to say exactly what that state [of the world] is. Let this state be S. Then, if a candidate system includes the Newtonian dynamical principles, one buys an enormous increase in the informativeness by adding a statement to the effect that at some time, the complete physical state of the

> world is S. For doing so will shrink the set of nomological possibilities to one. (Here I am taking advantage of the fact that Newtonian dynamics are two-way deterministic). But that is a disaster. (Hall (MS)).

One might hope that there *is* no such state. Unfortunately, Hall gives us a recipe for constructing one, by coding the position and momenta of each particle into a single constant. As position and momentum are natural properties, this satisfies the naturalness constraint (although, for some regularity theorists, it needn't). If this example is not convincing, other examples are easy to imagine: for example, a statement specifying the total energy of the universe, or a statement giving the total number of particles in the universe. Both of these would be highly informative in a deterministic system, but neither seems lawlike.

Real world examples are not hard to come by either. Take, for example, the flatness of the universe, or its smoothness. On large scales, every region of spacetime is approximately flat; similarly, on large scales, every region of spacetime has roughly the same amount of matter. The flatness and smoothness of the universe are not taken by cosmologists to be suitable candidates for lawhood; when these features of the universe were discovered, rather than being added to short list of fundamental laws, they were seen as surprising features in need of an explanation, and the current inflationary paradigm was partially motivated by the desire to explain them.

A similar point is made by Woodward (2013a: 8):

onlawful generalizations can be deduced from uncontroversial candidates for laws [...] in conjunction with appropriate information about initial conditions in [our system] S, and because we can hardly drop these uncontroversial laws from the best balancing systematization S*, our only alternative seems to be to exclude any information about initial conditions that might permit such derivations [...] the resulting system will not be strong.

Woodward presents this as a dilemma: our laws are designed to provide almost no information without information about initial conditions. One can include such information in the lawbook–then one arrives at a lawbook which is strong but contains lots of information which is intuitively contingent. Or one can leave it out–but then one will fail to satisfy the BSA's strength requirement. Woodward takes this to be evidence against the claim that strength, at least in Lewis's sense, is a feature which makes for lawfulness.

Earman and Roberts (2005a) concur while arguing that we should build our conception of the Humean base on our notion of boundary conditions. Unless our balance between strength and simplicity gives extreme weight to simplicity, importing some information about initial conditions will greatly improve our lawbook. But, as I'll now argue, no non-ad-hoc notion of simplicity can be found to save the BSA.

Laws and Scientific Practice

The regularity theorist should have as her goal encoding and explaining the practice of science. Our final regularity theory should both identify the criteria that scientists use to differentiate between laws and accidental regularities and explain why this distinction is such an important one. It's a mark against our theory if scientists, with full knowledge of their circumstances, would take a generalization to be a law despite its not meeting our philosophical criteria for lawhood. This mark is stronger if that generalization seems to do those things we take to be important for lawhood.

Regularity theorists recognize the tight connection between their views about laws and scientific practice: Loewer, in formulating his Package Deal Account, explicitly maintains that the laws should satisfy ``*whatever other conditions the scientific tradition places on a final theory*" in addition to maximizing strength and simplicity (Loewer 2007, emphasis in original). Hall eschews direct discussion of strength and simplicity in favor of imagining a ``logically omniscient perfect physicist" (LOPP) who weighs the virtues appealed to by scientists in determining which system is best (Hall MS). So evidence that scientists are not merely maximizing strength and simplicity is evidence against the BSA's criteria of maximizing simplicity and stregth.

Scientists are not willing to make trade-offs of simplicity for strength. Newtonian gravitational mechanics was preferred to Kepler's three laws not because of its simplicity, but because of its additional strength; general relativity supplanted Newtonian gravitational theory despite its considerable complications and variety of free parameters. Instead of sacrificing strength for simplicity, scientists start by formulating the strongest theories they can, given their evidence. If there are multiple such theories, they then choose from amongst these the simplest (this point is made forcefully by Woodward (2013a).

So simplicity considerations seem relevant in the choice between the Tychonic and the Copernican model, which are not observationally distinguishable, but seem misplaced in a choice between the Copernican and Newtonian model, as the Newtonian model makes strictly more predictions. Similarly, simplicity helps us determine whether our world has a Newtonian or neo-Newtonian spacetime; the latter is ontologically simpler, but (importantly!) just as informationally rich. When we employ Ockham's razor, we mean to cleave the fat from theories, but always leave their informational muscle.

Additionally, our interest in simplicity is not what one would expect, given the role of laws put forward by proponents of the BSA. Proponents of the BSA hold that we are interested in laws because we want as much information as we can wrap our heads around. Simplicity considerations, then, are intended to make the laws cognitively tractable.

But a vast literature in philosophy of science takes simplicity to not be desirable for cognitive tractability, but instead as an *epistemic* virtue. The justifications for invoking simplicity are many: some philosophers (Rosenkrantz (1977), Henderson (2013)) argue that rational priors give simpler theories higher credence. Others (Forster and Sober (1994)) hold that preferring simpler theories guards against measurement error.

Rosenkrantz (1977) explicitly formulates an objective Bayesian model of scientific inference, similar to one appealed to by Henderson (2013). In both cases, indifference principles are invoked: before we obtain evidence we ought, according to Henderson and Rosenkrantz, have equal confidence in all theories we believe are possible. We can think of this as a sort of fair-judgment axiom–if we start our inquiry with a higher credence in one theory than another, we will be biased. This will cause us to be unfairly resistant to evidence favoring the disapproved-of theory.

We can think of each theory as a family of equations–the family that agrees about everything except the values of its various free parameters. It is these families that Rosenkrantz and Henderson urge us to be indifferent between. We then divide the credence attributed to each family amongst its members, each of which sets the values of the free parameters. (I'll call the family of equations 'the theory family' and the theory with fixed parameters 'the theory'). What Rosenkrantz and Henderson then notice is that members of families with fewer free parameters thereby get a boost: because they're sharing their theory family's credence with fewer siblings, so to speak, they'll have higher prior probability than members of families with more free parameters. This higher initial credence will be amplified into a higher posterior credence after experiments are performed.

The takeaway is this: we prefer simpler theories not because they are cognitively more tractable, but merely as a foreseeable result of being unbiased about the form of our final theory before we accumulate evidence.

Statistical measures of simplicity don't rest on constraints on priors. Rather, they take simplicity to be a ward against measurement inaccuracy. If we choose a less-than-simple theory, we run the risk of failing to distinguish signal from noise. Simplicity considerations, on this view, are truth conducive: either they increase the ease with which we find the true theory (this is the value of the Bayesian information criterion) or they increase the likelihood that our next prediction is true (this is the value of the Akaike information criterion) (Forster and Sober (1994)).

In both cases, simplicity is valued not because it makes theories easier for us to understand, but instead because we are *more justified* in believing the simpler theory. For proponents of objective Bayesianism, this justification comes from *a priori* constraints on our prior probabilities. For proponents of various statistical measures of simplicity, this justification comes from the fact that favoring simpler theories is truth-conducive–either for our predictions or our belief in the relevant theory.

The history of science does not support the claim that, in searching for laws, scientists are looking for generalizations that best balance simplicity and strength. And our best normative accounts of simplicity considerations do not take simplicity to be related to the organization of knowledge. Taken together, these show that we do not and should not look for laws merely as organizational tools, as the BSA suggests.

An Illustrative Example

In §[boundary conditions] I argued that *our* world's best systematization contains statements which are not laws. In §[scientific practice] I argued that philosophical work on simplicity undercuts the idea that scientists weigh strength against simplicity. I'd like to supplement these arguments with another sort of counterexample to illustrate the fact that the BSA's simplicity and strength don't match our scientific and epistemic interests. In this world, there is no single best system, but, I argue, there is a set of generalization which play all the roles we need of laws. We should take this to show that the virtues identified by the BSA do not track the laws of the world.

TAYLOR'S WORLD: Take a world, T, which can be modeled by $F = ma$ together with some force laws. But at Taylor's world there is a true, informative statement about force which cannot be finitely stated. This force-statement is a function of a particle's velocity, and acts to damp the motion of the particle[8]:

$$f_T = -av - bv^2 - cv^3 - dv^4 \ldots$$

The faster something moves, the more Taylor's force opposes that motion. But each subsequent coefficient $(a, \ldots, d, \ldots)$ is much smaller than the one before it. So at low speeds it can be approximated as follows: $f = -av$, at moderate speeds approximated as $f = -av - bv^2$, and so on[9]. Scientists at this world are never able to fully specify this law, but active research programs engage in discovering better and better approximations and in determining which coefficients dominate at different speeds, enabling scientists at T to predict the results of any experiment with arbitrary, but never perfect, accuracy.

Because f_T is infinitely long, any systematization which contains it will be infinitely unsimple. So it will be no better with respect to the traditional measure of simplicity than a mere list of all facts[10]. Hence there will be no unique simplest, strongest systematization.

---

[8]This world is unlike ours in a variety of interesting ways, not the least of which is that it has a preferred velocity frame–some things are fundamentally at rest. But this shouldn't lead us to doubt that such a world is metaphysically possible.

[9]Taylor expansions like this one are used to model drag and friction, which are at our world understood to be nonfundamental, and to have no finite expression. We are merely imagining a world at which these laws are fundamental. But it's important to note that some scientific research at *our* world focuses on approximate laws like this, which are known to have no finite expression. These approximations seem to have many features of laws: supporting predictions and counterfactuals, being legitimate targets of scientific inquiry, etc. But this is not explicable on the traditional BSA.

[10]Because the force statement has infinitely many parameters, it will also be infinitely unsimple on measures of simplicity which depend on the number of free parameters a law statement contains. In conversation, Ned Hall has suggested that the law could be stated ``there are countably many coefficients $a_1, a_2, \ldots, a_n, \ldots$, such that there is a force $f_T = -a_1v - a_2v^2 - \ldots - a_nv^n - \ldots$" This would allow us to state our laws finitely, but would

According to the BSA, then, this would be a world without laws. But T is not an unordered world with too much chaos to allow coherent systematization. On the contrary! Science is as active at Taylor's world as it is here: controlled experiments can be performed, the findings of these experiments can be laid down in mathematical equations, which can then be projected to unobserved systems. Predictions and explanations relying on approximations of f_T can inform action, engineering, and research programs.

The situation of scientists at T illustrates two failures of the BSA. First, the BSA is unable to distinguish between lawbooks which have infinitely many free parameters. But more importantly, the BSA is blind to aspects of lawbooks that make science possible. A lawbook containing f_T supports scientific practice because it is applicable to isolated systems, conducive to approximation, and has free parameters which can be observed independently. A mere list facts has none of these features. So the simplicity requirement of the BSA is not capturing those features of laws which make science possible.

Purported Solutions

Most Humean solutions to these problems involve codifying the distinction between laws and sufficient conditions while denying that this distinction marks metaphysically distinct categories. Strategies of this sort are proposed by Hall (MS), Earman and Roberts (2005a),

---

significantly weaken the system. For example, this statement is plausibly true at the actual world, in the trivial case where all the coefficients are zero. But any statement which specified the values of the coefficients would again render the laws too complex.

and Woodward (2013b)[11]. As Hall's proposal is most explicitly a modification of the BSA, I'll focus on it.

Hall diagnoses the difficulty as follows: systematizations have two sorts of hypotheses, namely, *initial conditions hypotheses (ICH)* and *dynamic hypotheses* (DH). The ICH tells us which sets of initial conditions are nomologically possible; the DH tells us how those initial conditions evolve into others.

When we score a system for informativeness, we do not count all the laws equally. Rather, we give the systematization a higher score for having a more informative DH but a *lower* score if it has a more informative ICH: that is, the less we assume about the initial conditions the better. Hall correctly points out that this seems to match the motivation of practicing physicists and cosmologists, who regard a theory which requires special or specific initial conditions as less informative than one which does not require special initial conditions.

Although Hall's proposal is on the right track, it fails for two reasons, both of which I will examine in detail. The boundary between laws and initial conditions is not as firm as Hall's proposal requires. Although laws are typically dynamic, there are a number of historical examples of laws which contain static information or seem to constrain boundary

[11]It's not clear whether Woodward's view, according to which laws are invariant generalizations, counts as regularity theory. While Woodward expresses sympathy for Humeanism (which, recall, implies that if there are laws then some form of the regularity theory is true), his discussion of invariance makes it clear that invariance is characterized counterfactually, which makes him (in my terms) a modalist. But Woodward does advocate measuring strength by the inferential boost the laws provide: rather than counting a lawbook as stronger if it implies more, Woodward suggests that we count lawbooks as stronger when they, taken together with some particular facts, imply more particular facts. The view I defend is similar to Woodward's, but without counterfactual commitments.

conditions. And information which practicing scientists regard as part of the initial conditions can easily be woven into dynamic laws that make up the DH.

Secondly, as Hall admits, even if this proposal were actionable, this distinction cannot be motivated by the orthodox account of the role of laws as an organizational tool. I'll take these in turn.

Many laws contain, sometimes exclusively, nondynamical information. And the laws scientists give can easily be modified so that they include information we regard as parts of the boundary conditions of the universe. First, I take it that the paradigm case of a dynamic law-statement is a differential equation, in which at least one of the derivatives is with respect to time. This sort of law tells us how some properties of the universe, or a subsystem thereof, develop over time. Currently there are two such central equations in physics: Schrödinger's equation and Einstein's General Field equation. I'll focus on Schödinger's equation.

Schrödinger's equation contains free parameters which, when set, convey illicit information about boundary conditions. Schrödinger's equation is:

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

The Schrödinger equation makes use of the Hamiltonian operator, $\hat{H}$. Classically, the Hamiltonian represents the combined kinetic and potential energy of the system. Although things are slightly more complicated in the quantum domain, the Hamiltonian will still contain terms corresponding to the kinetic and potential energy of a system. But fully specifying the potential energy of a system amounts setting the initial conditions of that system; if the world contains no sourceless potentials (emanating in from infinity), we can express the potential energy of the system with a term for each particle.

Similarly, the wavefunction Ψ is a vector in a high-dimensional configuration space. This space has a dimension for the position and momentum of each particle. So by specifying the dimensions of configuration space we have thereby specified the number of particles in the universe.

But the number of particles in the universe is precisely the sort of information that is supposed to be a boundary condition, rather than governed by law–and certainly not by dynamic law. It's important to note that not every parameter in the Schrödinger equation intuitively is set by boundary conditions: $\hbar$ is taken to be a fundamental constant relating a system's wave properties to its energy. There is no important *syntactic* difference between these parameters[12].

The moral here is simple: fundamental dynamical laws contain free parameters which, when fixed, provide illicit information about the initial conditions of the universe. But they also contain free parameters which seem to provide no information about the initial conditions of the universe, and are suitably regarded as nomically contingent. The question then, is this: when is a parameter in a fundamental equation part of the dynamic law itself,

[12]The Hamiltonian is an operator rather than a constant. But other parameters have the structure of constants but are taken to be part of the initial conditions, such as the rest mass or charge of the particle being modeled.

and when is it a free parameter to be set by the boundary conditions? Neither the traditional BSA nor Hall's revision provide us with the tools for making this distinction.[13]

Here I've assumed that the Dynamic Hypothesis is a set of differential equations, and that the parameters of these equations can be fixed: for example, on my view, if $\mathrm{F} = \frac{\mathrm{dp}}{\mathrm{dt}}$ is a candidate law, then so is $\mathrm{F} = 4\mathrm{kg}\frac{\mathrm{dv}}{\mathrm{dt}}$, where the latter fixes the mass of every object in the world at 4kg as a matter of law. But this could be denied; we could instead hold that every free parameter in our dynamic law is part of the initial conditions of the universe. This is very implausible: some of the parameters of our physical laws, such as Coulomb's constant k or Planck's constant $\hbar$ fixed in the dynamic law in which they feature, and others are set by the initial conditions. These have their values with the force of nomic necessity. If our theory of laws requires a distinction between initial conditions and dynamic laws, we should have a clear and principled way of drawing this distinction which rules correctly in these cases[14].

Our second worry for this view is that not all boundary conditions are static and that not every law is dynamic. Physics is replete with nondynamic laws. Take, for example, Newton's Law of Universal Gravitation. This statement can be applied to any number of systems with different boundary conditions. But it is not a dynamic law. It tells us that the

[13]Philip Kitcher's (1989) unificationist model of explanation builds this distinction in on the ground floor (although Kitcher is not focussing on laws). Without getting bogged down in details, Kitcher takes the laws to be a set of equations with 'filling instructions', which tell us how to set each parameter. The informativeness of the lawbook is measured partially by how restrictive these instructions are. Kitcher's proposal is similar to mine, but it's worth noting here that nothing about the *syntax* of an equation distinguishes which parameter is set by initial conditions and which by law.

[14]Thanks to Ned Hall for discussion of this point.

gravitational force between two objects at some particular time is a function of their distances and masses, but not how the system evolves over time. Similarly, the second of Maxwell's equations, $\nabla \times \mathrm{B} = 0$, holds if and only if there are no magnetic monopoles. While lawlike, this is no more dynamic information than a law ruling out unicorns.

It's possible that a distinction between boundary parameters and nomic parameters can be made; and it's possible that a distinction between nondynamic laws and static boundary conditions can be made. But this cannot be done with the resources of the traditional BSA.

Even if this distinction could be made, I see no reason to do so. Without altering the BSA's theory of the role of laws, there is neither support nor guidance for any attempt to differentiate between the implications of a proposed element of our system, or to cordon of the dynamic from the nondynamic generalizations. Recall that the Mill-Ramsey-Lewis account does not merely purloin its scientific virtues from the practice of scientists– although hopefully it is inspired by that practice. Rather, it provides a reason for differentiating laws from accidental regularities and motivates its account of scientific virtues by appealing to this reasoned distinction. The role laws play, according to the MRL, is that of organizing our knowledge into a deductive system. It is easy to see how simplicity and informativeness are virtues given that purpose.

It is impossible to see how the distinction between dynamical and nondynamical information could in any way affect the degree to which an axiomatization organizes knowledge. The Best Systems proponent does not merely want to organize *dynamic* knowledge–she wants to organize *all* knowledge. Why then should she care whether organizing that knowledge requires her to make a highly specific statement about the world at some time?

As we've seen, *failing* to make *some* distinction is fatal to the BSA. There are clear cases in which scientists distinguish between boundary conditions and laws, so failing to do so removes the BSA's claim to scientific plausibility. But this distinction cannot be made by dividing the sentences of a lawbook into two camps. Providing a clear formulation of this difference and a believable motivation of recognizing it requires us to alter more than the letter of the Best Systems Account. We must uproot and rebuild it.

Laws and Epistemic Roles

The BSA is inadequate: it has too few virtues and lacks the features necessary to explain many features of scientific practice. The BSA should fall but the regularity theory of laws does not fall with it. In this section, I'll outline a new approach to the regularity theory. In §[The Epistemic Criterion] I argue that the BSA focuses on the output of scientific inquiry–predictions and explanations–and ignores the inputs–experimentation. In §[ERA] I develop an account of laws which takes them to be the midpoint of inquiry, resting between induction and prediction. In §[virtues] I make some suggestions about how this view accounts for scientific virtues, and in §[application] I show how this view handles the problem cases outlined in §[not good enough].

The view I arrive at takes the role of laws to be the primary metaphysical determiner of lawhood. This allows us to identify virtues which our lawbook should jointly maximize. So we can explicate the view in a best systems format, in terms of maximizing a set of scientific virtues. But I take the account primarily to be based in the role of laws, and talk of virtues to be essentially heuristic. This is because, as Baumann (2005) Callendar and Cohen (2009) and Woodward (2013a) justly complain, the best systems account give us no guidance in weighing the incommensurable virtues it identifies. The Epistemic Role Account put forward in §[ERA], by telling us for what purpose we are constructing our systematization, gives us a goal in terms of which we may judge tradeoffs between virtues.

The Epistemic Criterion

According to the BSA, the regularities that are laws occupy a special place because they help us organize our knowledge into a utilizable axiomatic scheme. This account of the role of laws leads naturally to simplicity and strength as unique scientific virtues. The unifying theme of the difficulties presented in §[not good enough] is this: the scientific method does not aim merely at organizing and unifying all truths. It aims at discovering truths that can be employed in a wide range of situations much smaller than the universe as a whole and at marshaling empirical evidence to provide epistemic support for believing those truths.

Fortunately, the BSA's account of the role of laws is not the only regularity account on offer. A competing account which gives the evidence-generating activity of science prime place in defining the laws is almost as old, but–with some justification–relatively overlooked. I'll call this account the *Naïve Epistemic Account*. The Naïve Epistemic Account holds that laws are those generalizations for which we have a high degree of inductive support. This account, like the BSA, can trace its roots to J. S. Mill's *A System of Logic*:

> These various uniformities, when ascertained by what is regarded as a sufficient induction, we call, in common parlance, laws of nature. (1843:187)[15]

According to the Naïve Epistemic Account, *R* is a law if and only if *R* has a high degree of inductive support. The NEA has never enjoyed strong support[16], and it's easy to see why. Here's a clear takedown by Fred Dretske:

> Laws do not begin to be laws only when we first become aware of them, when the relevant hypotheses become well established, when there is public endorsement by the relevant scientific community. The laws of nature are the same today as they were one thousand years ago. (Dretske, 1977).

The NEA is simply a nonstarter.[17] But at its heart is a kernel of truth. This kernel is also well-characterized by Dretske:

> Though laws are not merely well established general truths there is a related point that deserves mention: laws are the *sort* of thing that can become well established prior to an exhaustive enumeration of the instances to which they apply. (Dretske, 1977).

Dretske points out, accurately, that scientists are not merely looking for statements which, once known, underwrite counterfactuals, permit prediction, and enable us to give

---

[15]Mill provides these two incompatible characterizations of the laws of nature within a few pages of one another; it's not clear whether he recognized their incompatibility, or whether he intended either of them to be necessary and sufficient for lawhood. One is reminded of Hume's distinct and incompatible characterizations of *cause* in the *Enquiry*.

[16]It's not clear which philosophers have held the NEA. It's mentioned both by Ramsey (1927) and Goodman (1955), though neither accepts NEA or cites specific supporters. Whether Mill accepted the NEA or was merely describing inductive practice is an exegetical matter on which I have no stance.

[17]The NEA, distressingly, doesn't even require laws to be true. This makes it a mystery, given the NEA, why scientists would bother to advance their disciplines once they had some generalizations with a high degree of support.

explanations and perform manipulations. Scientists are looking for generalizations that can be known by observing a subset of their instances in controlled situations.[18]

Dretske is a modalist who believes that this second epistemic criterion can only be successfully accomplished by a metaphysically heavy account of lawhood. I disagree: the regularity theorist can identify features of generalizations which makes them uniquely positioned to be the target (non-exhaustive) inductive support.[19]

The ERA

The proponent of the BSA has in mind a scientist operating outside the universe and looking in. This ideal scientist starts with knowledge of all the facts of the world, so the only task left to her is to organize them. This idealization obscures a central aspect of scientific investigation. Even at the end of inquiry, when all truths have been discovered, the scientist will still have two jobs: organizing true beliefs *and providing evidence for them*. We do not merely want to organize truths; we want to organize *knowledge*. This requires the scientists to look for generalizations for which she can provide evidence.

The role of laws is not merely to support explanations, predictions, and counterfactuals. Laws also must be those truths discoverable through observation. Many of the most telling counterexamples to the BSA–such as generalizations stating the number of particles in the

---

[18]This is also noted by Hoefer (2007), for whom simplicity partially user-friendliness. ``User-friendliness is a combination of two factors: *utility* for epistemically- and ability-limited agents such as ourselves, and *confirmability*" (Hoefer 2007: 463).

[19]It's also worth noting that that the modalist, like Dretske, is no better at explaining why inductive practice is epistemically warranted than the regularity theorist. For a thorough discussion see Beebee (2011).

universe, its energy, or its exact state at some time–are intuitively not laws because they could never be supported experimentally. Noting this provides support for the *Epistemic Role Account of Lawhood*:

EPISTEMIC ROLE ACCOUNT (ERA): The laws of nature are those true statements which, as a group, are best suited to produce predictions and explanations and to be inferred from repeated observation and controlled experiments.

The ERA identifies a role for laws distinct from that provided by the BSA. This role includes both the *outputs* of and the *inputs* to science. The output-role that the ERA identifies is similar to that of the BSA: science should output a set of generalizations which will enable us to easily deduce predictions and provide explanations. Consequently we should expect strength and simplicity, or something like them, to be scientific virtues by the lights of the ERA. However, there is a slight difference: as we are looking for laws which produce predictions and explanations, we are looking for laws which provide a special type of information. This will lead us to modify the BSA's account of strength and fit.

The input-role of laws–that they must be suited to be inferred by observation and experimentation–does not appear in the BSA. Its inclusion will give us a tool for distinguishing boundary conditions from laws, and will force us to rethink the BSA's account of simplicity. It will also give us the resources to introduce a new set of scientific virtues which weigh against strength.

The requirement that the lawbook be supportable by observation or experiments, then, constrains our lawbook as follows: to perform experiments, we need laws which can be observed in isolated subsystems of the universe. And the laws must be observable in isolation. These are different requirements. The first requires the laws to apply to subsystems of the universe as well as the universe as a whole. The second requires the laws to be independently observable–parts of the lawbook must be observable while the action of others is minimized.

Science seeks to extend our knowledge from those contexts wherein we gain evidence to those in which we make predictions. Scientific laws have a central role to play in this extension: they occupy a place between induction (where we gain evidence) and deduction (where we apply it). Consequently our scientific system should allow us to identify quasi-enclosed systems where it can be applied. Then, it should tell us what features of these systems we should expect to be different between systems, and which features we should expect to be the same. The former are the boundary conditions, the latter the laws.

Before embarking on an ERA-based account of scientific virtues, I'd like to address a worry: the ERA's reliance on *observation* and *controlled experiments* might seem to make it problematically subjective or species dependent. It is not. The notion of *observation* in the ERA should be understood broadly, such that observations aided by complex, theory dependent apparati are suitable inputs for laws. Though I can easily imagine an alien species with novel sensory apparati, I cannot imagine an alien species able to make observations which we could not even in principle reproduce with the aid of some new technology[20]. I do, however, have serious difficulty imagining science proceeding without any sort of experimental input. This is what makes natural laws scientific laws, rather than mathematical or metaphysical truths.

[20]Perhaps that technology would involve th aliens themselves, either by incorporating them in some infernal machine, or by politely asking them

Scientific Virtues

These constraints motivate corresponding virtues, which the best candidate lawbook of our world must have. First: lawsystems are better if they provide information about subregions of the universe rather than about the universe as whole. This is not only necessary for us to gather evidence for them experimentally; it also allows us to make more predictions and provide better explanations, both of which typically operate locally rather than globally. Second, law systems are better if they have more *independently manipulable* parts. A similar requirement arises for chancy lawbooks: chance systems must be evaluated in terms of long-run frequencies of repeatable chance setups, rather than via the chance the system gives to the world as a whole (I'll discuss this idea more in ch. 3).

Each modification centrally makes use of the notion of an quasi-isolated subsystem of the universe. It is by attending to these subsystems that we bring the laws down to the realm of embedded agents. *Isolation* is arguably a law-dependent notion–a system is isolated if and only if outside influences are minimized, where outside influences are presumably characterized by the laws. Consequently I will briefly sketch an account of what I take a quasi-isolated subsystem–hereafter 'QIS'–to be[21].

[21]This notion is inspired by Cartwright (1999).

A subsystem is a QIS of the laws if and only if the laws are true of that subsystem; a lawbook is *true of* a subsystem if and only if the laws are true when any free parameters of that lawbook are filled by all and only those objects within the subsystem[22].

A QIS of the laws is a subsystem *described* by the laws; it behaves in accordance with the laws in the same way that the universe as a whole does. We can also appeal to an *approximate* QIS of the laws–this is a subsystem that the laws are *almost* true of.[23]

When introducing a notion which will play a central role in a theory, it's worthwhile to provide some examples: our solar system is a close approximate QIS of general relativity; if the variables of GR are filled just with all and only the objects of the solar system, the result is a true–or nearly true–sentence. The solar system is a less close QIS of Newtonian gravity. A particle accelerator is a QIS of high-energy quantum field theory.

Strength and the ERA

There are two ways in which a lawbook can provide information about subsystems of the universe: it can can provide information about more QISs, and it can provide more detailed information about each QIS. That is, a lawbook can model more, and it can model better. I'll call the former 'breadth' and the latter 'strength', as it's closer to the Best System's measure

[22]This does not amount to the requirement that the laws be generalizations. The laws could include reference to specific objects; if they did, then only those subsystems containing those objects would be QISs of the laws.

[23]Spelling out precisely what 'almost true' should mean here would take us too far afield. But the notion is not problematic. The laws of our world are given in terms of fundamental quantitative properties. These properties admit of a natural measure–1kg of mass is closer to 2kg than it is 20kg. A subsystem is an *approximate* QIS of the laws, then, if the values of the quantities in that subsystem do not diverge very far from those required by the fundamental equations. How far is too far? This is vague, but not more vague than the notion of an approximate QIS.

of informativeness–although importantly, unlike the BSA's notion of strength, this virtue is a function of local laws.

Both strength and breadth are valuable to our epistemic goals: a lawbook with more QISs has provides more opportunities for confirmation; a lawbook which provides more information about each QIS enables us to make more precise predictions. But they weigh against one another: as each subsystem will be qualitatively distinct, lawbooks can only increase their range of application by providing less precise information about each QIS.

Precisely formulating breadth is simple:

BREADTH: Lawbook L is *broader* than lawbook L* if and only if L has more QISs and approximate QISs than L*.

It's important to note that our interest in broad laws goes beyond just wanting laws that apply to a lot of systems; we want laws that apply to a lot of *different* systems. We are interested in rules that we can take with us from our observational and experimental contexts and apply to situations quite qualitatively different from them. So the breadth of the lawbook should be counted not just by adding up all of the QISs to which the laws apply, but instead by focussing on those lawbooks that apply to many quite qualitatively different systems.

A broader lawbook allows us to observe the laws in action in more situations. General Relativity is an extraordinarily broad set of laws: its approximate QISs include every star system, galaxy, and galaxy cluster. Similarly, quantum field theory is excessively broad: its QISs include every nearly isolated atom or molecule. Of course, not every subsystem of the universe (or solar system) is a QIS of these laws: the Pluto-Sun system is not a QIS of GR, as other objects in the Kuiper belt exert a strong gravitational influence on Pluto.

Precisely formulating a local virtue of strength is slightly more complicated. Recall Lewis' account of strength for the best system:

GLOBAL STRENGTH: Lawbook L is *stronger* than lawbook L* if and only if L rules out more worlds than L*.

A *globally strong* lawbook rules out more possibilities than a globally weak lawbook. But it may rule out all but a small set of worlds without giving much specific information about the internal structure of the world; it may provide only global information (such as the total number of particles), information which cannot turn up in any models except the one big model–the universe as a whole. Lewis' notion of informativeness is irredeemably global. Modifying GLOBAL STRENGTH to insure that the lawbook gives us the sort of information we need to meet our epistemic goals requires us to formulate a local measure of informativeness. To do this I'll appeal to the notion of a *counterpart* of a QIS. A world w is a counterpart of a QIS q if and only if w contains all and only counterparts of the objects in q, and the properties of and relations between the objects in w correspond to the intrinsic properties and relations of objects in q. The solar system is a QIS of general relativity; its counterparts are worlds which contain counterparts of all and only those 8 planets and many more asteroids and meteors with in the gravitational influence of the sun. We can now provide a local measure of informativeness:

STRENGTH: A lawbook L is globally stronger than lawbook L* if and only if L rules out more counterparts of its QISs (and approximate QISs) than L* does of its QISs (and approximate QISs).

Local strength can only come at the cost of breadth. Increasing the breadth of a lawbook requires us to make the laws compatible with more actual subsystems; but unless the world's subsystems are highly uniform, this will require the laws to be compatible with more QIS counterparts. And so the laws will be less strong.

The account of strength offered here should be seen as a first-blush account of strength. For we don't count all information equally when we search for laws; some information is cheap and easy for us to obtain when we encounter a new situation. The laws are most useful when they can take this cheap information and transmute it into something more valuable. So we look for laws that amplify or expand on the information we already have about a QIS, rather than giving us information on their own. Rather than discuss this at length here, I'll take it up in more detail in ch. 2.

Simplicity in the ERA

The view we've sketched so far has three virtues: strength, breadth, and fit. On the traditional BSA, the need for strength and fit is reined in by the requirement that the laws be simple. But as we've seen, *breadth* plays this role in the ERA: if laws provide too much information about their QISs, they have fewer QISs. This lowers their breadth, and makes them harder to discover through empirical investigation. What need, then, is there for simplicity?

The answer can be found in some of the criticisms laid against the BSA's motivation of simplicity. Recall from §[scientific practice] that most philosophical accounts of simplicity take simplicity to be a virtue not because simpler laws make for better axiomatizations of knowledge, but instead because simpler laws enable induction to proceed more smoothly. The three accounts we considered–objective bayesianism and the akaike and bayesian solutions to the curve-fitting problem–take favoring simplicity to be required by *a priori* constraints on rationality, the fastest route to the correct model, or the best way to minimize predictive error, respectively. In each case, the syntactic simplicity is show to combine with independently motivated constraints on inductive reasoning encourage us to favor simpler theories.

Simpler theories, then, are favored by standard inductive practice–though, plausibly, not at the expense of strength. Nonetheless we can add a simplicity requirement to the ERA's list of scientific virtues, with a slight tweak: both solutions to the curve-fitting problem and objective Bayesian accounts of simplicity take the simplicity of a theory to be a function of its free parameters, not its syntactic length. This should be reflected in our theory of laws.

SIMPLICITY: Lawbook L is simpler than lawbook L* if and only if L has fewer free parameters than L*.

Modularity

Our lawbooks are not given to us all at once; rather, we must piece them together a bit at a time. Thus it was nearly one hundred years after Newton's three laws, together with his Law of Universal Gravitation (1687), that the final ingredient of Newtonian Mechanics was added to the classical lawbook with Coulomb's law (1785). And classical mechanics didn't achieve its final form for *another* century, when Maxwell provided a unified theory of electromagnetism.

This piecemeal method of scientific discovery is matched by a divide-and-conquer methodology of evidence gathering. Each part of the lawbook must be independently tested; every fundamental constant must be observed in isolation. Our discovery of Newton's three force laws relied crucially on the existence of QISs, like the solar system, in which only one force dominated. The fundamental constants which determine the relative values of the fundamental forces can only be ascertained if each force can be observed independent of the others

Thus to discover and test our lawbook we need subsystems that are QISs of only a some of the laws. I call this virtue *modularity*. In order to explicitly define modularity I'll need a notion of a portion of the lawbook.

REDUCTION: A lawbook *l* is a *reduction* of lawbook L if and only if (a)*l* contains a subset of the laws of L, or (b) some of the free parameters of L are zero or held constant in the laws of *l*.

So the solar system is a QIS of a reduction of the laws of classical mechanics: it can be described without Coulombs law, or with all charges (a free parameter of Coulomb's law) set to zero. It is this feature of the solar system which allows us to observe the action of the law of universal gravitation isolated from other parts of the lawbook. The laws of classical mechanics are *modular*.

MODULARITY: Lawbook L is more *modular* than lawbook L* if and only if there are more QISs of *reductions of* L than of *reductions of* L*.

It's worth noting that modularity trades off against simplicity. The more free parameters a lawbook has, the more reductions it admits, and so the more QISs of reductions it can have. This should not worry us–modularity only encourages us to increase the complexity of the laws when doing so enables us to independently test our assumptions, or adds lower rungs to the ladder of scientific discovery.

Applying the ERA

The Epistemic Role Account of natural law takes laws to be those generalizations that are best positioned to sit between inductive learning and deductive predictions and explanation. I've argued that such laws will balance *breadth*, *strength*, *simplicity*, *fit* and *modularity*. Consequently I tenuously characterize the ERA as follows:

ERA $_{\text{virtues}}$: The laws of nature are those true generalizations that best balance breadth, strength, simplicity, fit, and modularity.

How, then, does the ERA circumvent the problems with the BSA? Recall that our problems with the BSA were:

1. The BSA does not provide us with the resources to distinguish initial conditions from laws.
    1. Many laws are static.
    2. Laws contain parameters, some of which are regarded as universal constants and other of which are set by the boundary conditions of the system being studied.
2. The BSA does not give us any guidance as to how its apparently incommensurate virtues weigh against one another.
    1. Scientists do not weigh strength against simplicity.
3. The BSA does not allow for laws with infinitely many free parameters.

The ERA successfully overcomes these difficulties.

According to the ERA, the distinction between initial conditions and laws is not a distinction between generalizations and particular statements, nor is it a difference between dynamic and static equations. It is not a syntactic distinction at all; instead, whether some truth belongs amongst the laws or the initial conditions depends on how many subsystems of the universe it's true of. Those equations which hold in many isolated subsystems are laws; those which do not, no matter how *globally* informative they are, sit with the boundary conditions. Force laws, though static, hold in many subsystems; constants like Planck's and Coulomb's are the same in all subsystems; but systems vary widely in their energy and number of particles. So the former are part of the laws, and the latter part of the initial conditions.

Many authors, especially Callender and Cohen (2009), have justifiably complained that the traditional BSA appeals to a best balancing of scientific virtues, but gives us no guide as to how we balance them. The BSA's virtues of simplicity and strength seem incommensurable, and as Woodward (2013a) points out, we cannot look to scientific practice for help, as scientists seem never willing to trade in their strength for a gain in simplicity.

The ERA has more resources. Because the virtues are motivated pragmatically, by their connection to the epistemic role of laws, we can appeal to the role of laws to determine which balance is best. When are we willing to give up strength? When sacrificing breadth would leave the laws too narrowly applicable to be discovered or tested. When does simplicity favor one putative lawbook over another? When independently motivated constraints on induction would draw us to the first lawbook rather than the second (this gives us little motivation to sacrifice strength for simplicity, but explains why we favor a simpler lawbook over a more complex, but equally strong, lawbook). How modular must the laws be? Modular enough for us to discover the fundamental constants, and to bootstrap our way into discovering the whole book.

The BSA was unable to show why Taylor's force law is intuitively lawlike. Because it has infinitely many free parameters, the BSA gave it the same score as a list of all facts at that world. The ERA places more weight on the sort of information conveyed by the laws than their syntactic features. Because a lawbook containing just Newton's laws and a force law–no matter how complex–has more QISs, the ERA correctly rules that its statements are laws, and a list of facts is not. Such a lawbook provides the *sort* of information we want: it is both locally stronger and broader than a lawbook which merely lists facts.

I conclude that the ERA meets the criticisms of the BSA.

Conclusion

The Best Systems Account of laws was designed to distinguish laws from merely accidental generalizations. I argued in §[BSA] that the BSA was tied to a picture of science as an enterprise focused on organizing knowledge. This account of laws and of the role of laws in science has met with insuperable difficulties, as I argued in §[not good enough].

But the regularity theory's picture of laws as mere generalizations that play an important role has not met any difficulties it cannot overcome. The letter of the Mill-Ramsey-Lewis view must change, but its spirit is preserved in a less naïve view, the Epistemic Role Account of laws. The Epistemic Role Account of laws can overcome the counterexamples to the BSA without giving up its metaphysical scruples.

The following chapters seek to extend these ideas to explain other features of our scientific worldview. In ch. 2 I show how this view extends to nonfundamental laws; as I will argue, nonfundamental laws provide information that is useful to us but absent even in the most informative set of fundamental laws. In ch. 3, I show how this view can be extended to solve some problems plaguing the traditional BSA's account of fit.[24]

[24]Thanks to Thomas Blanchard, Marco Dees, Heather Demarest, Ned Hall, Barry Loewer, and Jonathan Schaffer for discussion and comments on drafts of this chapter.

Coordinating the Sciences

Introduction

The view of laws just presented is meant as a view of fundamental laws; probably, our most fundamental science is physics. But not all of our laws are fundamental laws. Laws and counterfactual supporting generalizations also appear in biology, chemistry, economics, and psychology. In this chapter, I will show how the view of laws just outlined connects to the laws in these sciences.

A combination of independence and mutual constraint characterizes the relationship between the sciences. All sciences are able to employ the same methodology as physics, in terms of developing a conceptual structure, using that structure to formulate explanations through laws, and holding those concepts and laws accountable to the world through experimentation. The generalizations they arrive at support counterfactuals and feature in explanations [25]. Despite this independence, the sciences exercise mutual constraint on one

---

[25]Woodward (2003), chapter 6, denies that the generalizations of the special sciences are laws; in doing so, he rejects the notion that only laws are counterfactually invariant, and that only laws are available for use in explanations. Woodward's reasons are simple: according to standard accounts of law, laws must be exceptionless. But the generalizations which feature in special scientific explanation are not spaciotemporally unrestricted and have exceptions. For an argument that there are no laws in biology, see Beatty (1995); for a response, see Mitchell (2000).

Like Woodward, I have no truck with a verbal dispute about the word 'law'. Here, and throughout this paper, I will use 'law' to refer to those counterfactually robust generalizations that can underwrite predictions and feature in explanations. Claiming that the generalizations of the special sciences are not laws will not remove the burden of another: even counterfactual disagreements between sciences show that at least one set of laws contains a falsehood. And these sciences exhibit a hierarchical explanatory structure: interscientific explanations flow up, from more to less fundamental sciences. Accounting for these four features is the job of a philosophical account of law in the special sciences.

This problem has generally been approached as the problem of *reduction:* which sciences *reduce* to which? Specifically, do all sciences reduce to physics? And how is the relationship of reduction to be understood? Understanding the problem as a problem of reduction is mistaken for two reasons: firstly, it biases the discussion against views which emphasize the methodological independence of the sciences. Secondly, it creates the illusion that we are looking for a simple yes-or-no answer. Disagreements over whether, e.g., mere supervenience is sufficient for reduction distracts us from the underlying features of the relationship between sciences that need to be explained. To avoid these confusions, I'll call the puzzle posed by the relationship between the sciences the *coordination problem*: how are various scientific disciplines coordinated with one another?

In this paper, I'll present a new solution to the problem of coordination. But first, I'll identify two strains among extant solutions to the problem of coordination. I call these the *imperialist* and the *anarchist* solutions to the coordination problem. The imperialist sees the special sciences as a consequence of fundamental physics; the laws of the special sciences are laws because they can be *derived from* or *grounded* in the laws of physics. This strong reductionist view seeks to make every explanation an explanation from physics. The

---

explaining these features of those generalizations, and so will not (by itself) solve the coordination problem. Consequently I will not address the question of whether laws must be spaciotemporally unrestricted.

anarchist, on the other hand, denies that the sciences are connected. Rather, she sees them as each unifying a body of facts, or cataloging the dispositions of properties, with no connection to any other science.

Both of these views fail to solve the coordination problem; the imperialist fails to account for the independence of the sciences, and the anarchist fails to account for their mutual, asymmetric dependence. I'll conclude by offering a third view, which I call *the democratic view*: on my view, the various sciences work together to generate a set of laws, the informativeness of which are evaluated holistically. But because various scientific disciplines are epistemically isolated, in a way in which I will make more precise, they add to this lawbook semi-autonomously. The view I advocate has the advantages of both the imperialist and the anarchist view. Like the anarchist, and unlike the imperialist, I hold that the laws of the special sciences are *made laws* in the same way that the laws of fundamental science are. Like the imperialist, but not the anarchist, I hold that the laws of physics are fundamental, and that there is an asymmetry between the special sciences and physics.

Though the divide between imperialism and anarchism crosscuts views about the metaphysics of laws, the proposal I offer depends on features of the view of laws outlined in ch. 1. This is a Humean view: it relies on no fundametnal notions of necessity or dependence. But nonhumeans will find much here to like: Humeans take the epistemic role of laws to be *constitutive* of natural lawhood. That is, they believe that laws support counterfactuals and provide explanations because of their epistemic utility, not vice versa. A modalist about laws, who takes laws to have either irreducible nomic or metaphysical necessity, will still need to understand the epistemic utility of laws, and so can tack this account on to their more metaphysically robust account as an explication of the epistemology of laws. And many modalists about law take the only truly necessary laws to be those of fundamental physics; such a metaphysician of law can accept this view as an account of the laws of the special sciences while denying that it is a sufficient account of lawhood *simpliciter*. Finally, I argue in Section 1 that neither the imperialist nor the anarchist provide an adequate solution to the coordination problem. But thus far all modalist accounts of law fall into one of these two camps. So a modalist need either respond to the challenges presented in Section 1 or reject one of the four features of coordination there identified.

The Imperialist and the Anarchist

In what follows, I will first set out this dichotomy in broad strokes and then show how individual philosophers fit into one or another camp. It's worth noting that views about the relationship between physics and the special sciences crosscut views about the metaphysics of laws; although ultimately I favor a broadly Humean view of laws, my criticisms of the current theoretical space of possibilities do not rest on any metaphysical scruples.

To help illustrate the difference between the anarchist and the imperialist, and later to elucidate the democratic view, I'll make use of an idealized epistemic agent. She needs, unlike us, to have a vast capacity for absorbing and combining information from various sciences. But we will not assume–except when a view of laws demands it–that she is logically omniscient, or that she, like Laplace's demon, is able know everything about the state of the world (though she might), nor will we assume that inference is for her without computational costs. Some of these details of our agent will be fixed by the various purported solutions to the coordination problem. We can refer to her as a FISA: a Fairly Ideal Scientific Agent[26]. She will have a set of conditional credences, and these conditional credences will reflect encode the laws of various sciences.

If one of our laws says that *if A then B*, FISA's credence in B conditional on A will be 1. But the laws FISA responds to need not be deterministic: if our laws are statistical, this to will be reflected in her credences. So if it is a law that agents who are asked to memorize a ten-digit number are more likely to utter racial slurs than those who have no number to remember, her credence F(slur|number) will be less than her credence F(slur|~number)[27].

I will evaluate imperialism, anarchism, and democracy with respect to four features of the relationship between physics and the special sciences (briefly introduced in the introduction). These desiderata must be a bit vague; different views about the relationship between the sciences should be allowed to provide slightly different account of what, for example, the asymmetric dependence between physics and biology amounts to.

- METHODOLOGICAL INDEPENDENCE: Each science is able to formulate generalizations and support them evidentially via induction, and each science is able to determine its own conceptual structure.

[26]The strategy of explicating views of laws via an idealized scientist is becoming more common, and appears in Callender and Cohen (2010) and Hall (MS).

[27]Typically, discussions of objective probability assume that the objective probabilities are precise in situations in which they are defined. But this is not obviously the case for some special science generalizations: plausibly, some laws in the special science provide comparative relations between conditional probabilities without nailing those probabilities down. While I think that a complete account of special scientific law should be compatible with this (and believe that mine is), addressing this issue is beyond the scope of this paper.

- COUNTERFACTUAL ROBUSTNESS: The generalizations of the special sciences are counterfactually robust: that is, they both support counterfactuals and hold in a variety of counterfactual situations–including, plausibly, counterfactual situations in which the laws of lower-level sciences do not hold.

- MUTUAL CONSTRAINT: Distinct sciences cannot make inconsistent predictions, including predictions about what would occur in merely counterfactual situations, and cannot provide inconsistent constraints on belief or credence. Closely related sciences are such that the entities studies in one science can be located amongst the entities studied in another, often via a functional reduction.

- ASYMMETRY: Metaphysical or grounding explanations between sciences go in one direction only; this direction of explanation creates a hierarchy roughly lining up with the direction of mereological dependence, where the entities of higher-level sciences are made up of the entities of lower-level sciences. One way in which this asymmetry manifests itself is as follows: entities and behavior at the higher level can be located amongst the entities studied at the lower level; higher-level regularities are often targets for explanation at the lower level[28].

[28]An excellent example of this is the reduction of chemistry to physics, where chemical kinds–elements–are taken to be arrangements of physical kinds–protons and neutrons. The stability of some arrangements of protons and neutrons but not others explains the limited number of elements; the physical properties of these arrangments, such as the allowable energy levels of electrons orbiting them, explains the chemical properties of the elements in question, such as electronegativity.

We are looking for a view of laws that explains these four aspects of the relationship between the sciences while retaining descriptive adequacy: the closer the laws posited by the view resemble those of our current sciences, the better.

Our goal is to understand the practice of science: scientific discovery happens in diverse disciplines, using diverse methods. Our current sciences provide our best example of a working scientific hierarchy. So it should be believable that the solution under discussion is a view about the laws of our sciences–if the view does not allow some special scientific generalization to be a law, or requires us to add to the fundamental laws, this is a demerit of the view. This is a defeasible requirement. For the laws we have now are not the final laws; and the divisions we now carve between our sciences are somewhat arbitrary. So a philosopher has it within her rights to argue that our final theory will have features no current theory has; and she may likewise argue that some laws which are currently considered to be in one science actually belong in another–or that the division between two sciences isn't a division we should be worried about.

The Imperialist

The imperialist view holds that the lawhood of the laws of the special sciences derives from the lawhood of the fundamental laws[29](or the laws together with 'robust' initial conditions). The imperialist may hold that they can be *derived* from the fundamental laws,

[29]It's important to bear in mind here that we are discussing the dependence of *laws* on *laws*. Any physicalist philosopher is committed to some dependence of all higher-level facts, including the facts about which generalizations are laws, on the physical *facts*. But this dependence need not go directly through the laws; the unifying claim of imperialism is that the *lawhood* of the special sciences is dependent directly on the laws of physics (together with some other physical facts).

but she need not: she may hold instead that they are metaphysically necessitated by the fundamental laws, or that they are *grounded* in the fundamental laws–where A grounds B only if A metaphysical necessitates B and A explains B.

In this section I will outline imperialism. Then, I will present a dilemma for imperialist views. Either they are *austere*, and allow only a small set of fundamental laws to ground the lawhood of the special sciences, or they are *permissive*, and allow for a wide variety of fundamental facts to ground the lawhood of generalizations of the special sciences. On the first horn, austere views provide too little fundamental stuff to ground the wide variety of laws we find in the special sciences. On the second horn, permissive views both can neither ground the counterfactual robustness of the special sciences nor properly specify which fundamental facts are suitable to ground the special scientific laws.

A prototypical–if dated–imperialist is F. P. Ramsey (1927), who held that there are three grades of law: fundamental laws, laws that are derived from the fundamental laws alone, and laws which are derived from the fundamental laws and some 'robust' initial conditions. We might add a fourth category, not available to Ramsey: laws derived from the fundamental laws and *a posteriori* necessities, like 'water=H_2O'[30]. Finally, we should remember that it is open to imperialists may add to the set of fundamental laws so that they have sufficiently strong implications for the special sciences.

[30]Of course there's a fifth possible type of law: one dependent of the fundamental laws, robust initial conditions, and *a posteriori* necessities. But these will not improve the situation for the imperialist: I will argue that neither initial conditions nor *a posteriori* necessities can ground the laws. If these cannot solve the coordination problem on their own, neither can the two together.

To illustrate the imperialist view, let's look at the behavior of an idealized scientist–the FISA. Our FISA, according to the imperialist, starts with a set of fundamental laws. These laws may be sentences which together maximize strength and simplicity, as the Humean holds (Lewis (1980, 1983), Beebee (2000), Loewer (2007, 2008, 2009)), they may be generalizations which are backed by a relationship of necessitation between universals (Armstrong (1983, 1997)), or they may be sentences which describe the dispositional essences of the properties which feature in them (Ellis (2000, 2001), Bird (2007)) She then works out the consequences of these laws. On the most austere view, her conditional credences now encode the fundamental laws and the laws of the special sciences. But on more permissive views, she isn't done. On a permissive view, all she has now are the fundamental laws. She may still conditionalize on either some special set of the initial conditions, or she may conditionalize on *a posteriori* necessities–typically property identities. Once she has done this, says the permissive imperialist, she has at her disposal both the fundamental laws and those of the special sciences.

It's worth noting here that the laws of the special sciences need not receive probability 1. Indeed, likely they should not. For the laws of the special sciences are not exceptionless, as are the laws of physics. So an adequate account of special scientific law, imperialist, anarchist, or democratic, must do one of two things: Either it ought to hold that the conditional credences assigned to the special scientific laws are not unity, or it ought to provide some guidance to the situation in which the laws fail–a *ceteris parabis* condition.

Before we look at the problems with imperialism, we should note its advantages. Imperialism clearly and coherently explains two features of the coordination problem:

MUTUAL DEPENDENCE *and* ASYMMETRY. According to imperialism, the laws of the various scientific disciplines must be compatible because some of them are a consequence of others (together, for the permissive imperialist, with robust initial conditions or *a posteriori* necessities). If we discover a contradiction between the apparent predictions of two sciences, it's impossible that one of them is derived from the other. Consequently one of them must have the wrong laws[31]. And the ASYMMETRY of the sciences is neatly explained as well, because the laws of less fundamental sciences are a consequence of those of the more fundamental science, but not vice versa. The asymmetry of the sciences is just the asymmetry of deduction: the special scientific laws follow from those of physics, but not vice versa.

As to the METHODOLOGICAL INDEPENDENCE of the special sciences, the imperialist gets a weak pass. For the imperialist is not committed to our FISA actually representing scientific reasoning; we may not be able to perform the computations which FISA performs. She is fairly ideal, and so may be ideal in ways in which we are imperfect. So–perhaps–we with our limited cognitive resources are forced to engage in standard inductive reasoning to discover the laws of the special sciences, rather than simply deriving them from the laws of physics (together with whatever else). According to the imperialist, the fact that some

[31]Imperialism doesn't hold that the mistaken science must always be the special science; we might take physics and thermodynamics together to be fundamental, but recognize that the contradiction between physics+thermodynamics and geology, noticed by Kelvin in the 19th century, told against the then-dominant theory of physics rather than the then-dominant theory of geology. Because the geological laws yielded a different age for the earth than physics+thermodynamics, and because if B contradicts A, it's not the case that A implies B (given that A is self-consistent), we know that one of A or B must *not* be a law. We don't know whether to take this as a modus ponens of ∼B or a modus tollens of A.

special science generalization is inductively supported is strong evidence that it is a law, and so strong evidence that it is a consequence of the laws (and 'robust' facts) of physics.

This pass is a weak one. For the imperialist has given us no reason–at least not yet–to believe that the inductively supported generalizations of the special sciences will line up with those derivable from physics. Note that it is not enough for the imperialist to note the counterfactual robustness of special scientific laws and claim that this robustness must come from the laws of physics. For the source of this counterfactual robustness is precisely what is at issue[32]! Rather, she must provide some independent reason to believe that higher-level inductive reasoning will arrive at the consequences of physics, rather than some other generalizations.

Despite these successes imperialism lacks the resources to explain COUNTERFACTUAL ROBUSTNESS while retaining descriptive adequacy. To see this, let's first examine *austere imperialism.* Austere imperialism holds that the laws of the special sciences are a consequences of the laws of physics alone. We can see right away that austere imperialism will simply not do: for the laws of physics alone have too few direct consequences to underwrite all of the special science laws. And this is reflected in the structure of the laws of physics and the laws of the special sciences. The laws of physics are temporally symmetric, exceptionless, and deterministic[33]. The laws of the special sciences are

[32]Loewer (2008) argues that, because the higher-level frequencies are determined by statistical mechanical probabilities, observations of higher-level frequencies give us evidence about the underlying fundamental probabilities. I will address this later.

[33]Quantum mechanics, on either the orthodox or Ghirardi-Rimini-Weber formulation, is indeterministic and temporally asymmetric. But this should not concern us: first, the orthodox interpretation is widely regarded to be inadequate, both in specificity (it posits collapses, but does not say when or how they occur) and in internal consistency (the

temporally asymmetric, have exceptions, and are often statistical. So the special scientific laws could not be a result of the laws of physics on their own[34][35].

Now consider the permissive imperialist who adds *a posteriori* necessities. It's not at all clear how this could help. For if the laws of physics are temporally symmetric, exceptionless, and deterministic, adding a metaphysically necessary lasso between these laws and some higher-level terms will not introduce an asymmetry, exception, or indeterminism.

So to retain descriptive adequacy, the imperialist ought to become more permissive. She ought to include, not only the laws of physics and *a posteriori* necessities, but also some 'robust' initial conditions. To make this work, she will need a clear notion of robustness: one which will lead to an explanation of the lawhood of special scientific laws. By adding facts about the past, and not the future, we can secure the temporal asymmetry, exceptions, and indeterminism of the special sciences. But note that adding these initial conditions immediately makes this aspect of the coordination problem more pressing: for if the initial

---

indeterministic collapse postulate is in tension with the deterministic evolution of the wavefunction). Meanwhile, the GRW interpretation makes predictions which are distinct from those of orthodox quantum mechanics, but enjoy limited empirical support. In either case, it's doubtful that the temporal asymmetry and indeterministic nature of quantum mechanics underlies the asymmetry and indeterminism in the special sciences. Finally, on either of the other two leading interpretations of QM (Bohmianism and Everettianism), physics is deterministic and temporally symmetric.

[34]This is extremely clear if the dependence relation is something like derivability. But more permissive dependence relations, like supervenience or metaphysical grounding, face the same problem: if the supervenience base, or grounding facts, are temporally symmetric, exceptionless, and deterministic, how can they by themselves ground asymmetric indeterministic laws?

[35]For a more thorough and engaging discussion of this problem, see Loewer (2008).

conditions are not *themselves* laws, how can they make *other* generalizations laws? It seems that the imperialist must talk fast if she is to explain COUNTERFACTUAL ROBUSTNESS (it is just this issue that leads Beatty (1994) to argue that biology is without laws).

The contention here is not that accidental facts never support counterfactuals. They do: the accidental fact that my favorite mug just appeared on a TV show makes it the case that *if I were to sell it on Ebay, I would make $70.* The worry is instead that the laws of the special science are robust in a way that these accidents are not. The fact that all of the coins in my pocket are quarters makes some counterfactuals true: for example, *if I were to take a coin from my pocket, it would be a quarter*. But others are not supported: it's not the case that *if this nickel were in my pocket, it would be a quarter*. The laws of biology are not like this: it's true that if I were a bear, I would hibernate through the winter. This second class of counterfactuals, about what would occur under some manipulation, is the sort of counterfactual that can be grounded by laws but not accidents.

It's important to get this difference right; the reasons we look for laws require them to be robust under counterfactual situations. We want laws which will help explain features of our world, will enable us to predict what will happen if we act in different ways, and empower us to make decisions based on the causal structure of the world. To play these roles, laws must be sufficiently counterfactually robust to hold in situations in which we perform manipulations.

Next, without a specification of which initial conditions are *robust*, the imperialist's solution to the problem of METHODOLOGICAL INDEPENDENCE is even more fraught. For whichever initial conditions she chooses, she will need to explain why *those* initial conditions, and not the others, make a generalization available for inductive discovery at the higher level. But there is no reason to believe that there is any set of conditions on robustness that will do this. In fact, there is reason to believe the opposite.

The challenge for an imperialist is to find a set of facts which (a) together with the laws ground the laws of the special science, (b) do not mistake accidents for laws at the special scientific level, and (c) are sufficiently counterfactually robust to underlie the counterfactual robustness of the special scientific laws. In order to satisfy the third disederata, the facts appealed to by the imperialist must be in some sense unified; if they are not, the imperialist view will lack the resources to explain the counterfactual robustness of these generalizations without succombing to *ad hoc*kery. But these three conditions have yet to be met: imperialist theories either have too little in their grounding base, in which case they don't explain the lawhood of all special science laws, or they have too much in their grounding base, in which case they don't explain the accidental nature of special-scientific non-laws, or fail to identify a non-*ad-hoc* set of initial conditions to add to the base, in which case they are unable to explain the counterfactual robustness of the special sciences in terms of the counterfactual robustness of physics, in which case they have failed to deliver on their imperialist promise.

To see this, we may do well to examine one of the most worked out extant imperialist theories: that of Loewer (2008, 2009). Loewer recognizes that initial conditions on their own cannot a counterfactual support; so he argues that some initial conditions ought to be included in the book of laws. Specifically, he thinks that, in addition to the laws of physics, our fundamental lawbook should include PROB, "a law that specifies a probability distribution (or density) over possible initial conditions that assigns a value 1 to PH [the initial low entropy condition] and is uniform over those microstates that realize PH," (Loewer, 2008:19). As this low-entropy initial condition is a law, it is just as able to underwrite counterfactuals as the other laws in our fundamental lawbook. And PROB, Loewer argues convincingly, deserves to be in our lawbook for the same reason other laws are: adding it dramatically increases the informativeness of the lawbook without unduly complicating it.

So far, Loewer looks to have solved the problems of austere imperialism without adding the paralyzing complications of the permissive view. PROB is temporally asymmetric and probabilistic, and so can underwrite similar temporal asymmetries and probabilistic higher-level laws that don't follow from physics alone. But because PROB (according to Loewer and Albert) is a law, it neatly explains the counterfactual robustness of its consequences.

Unfortunately, PROB and the laws of physics cannot save imperialism. They are, by themselves, too permissive: many generalizations will have high probability, according to them–more than are counted as laws by the special sciences. This is because many highly probable generalizations will be burdensomely gruesome: we can take any two special scientific laws, which we can assume are given a high probability by the Loewer-Albert system. We can then define gruesome predicates by pasting together terms from each law, and thereby arrive at a gruesome generalization at least as probable as the conjunction of the two laws. If they have a high enough probability, this generalization will also have a probability above whatever threshhold we set for lawhood, but because of its gruesomeness, will not be a law.

And there is no guarantee that the actual laws of the special sciences will be given a high initial probability by these two. To see this, consider a law of population genetics. Such a law will depend sensitively on contingent facts early in the evolution of modern animals (it is just this problem which is discussed in Beatty (1995)). But PROB does not give a high probability to these historical facts–or at least does not probilify them over their alternatives. So it is unable to distinguish the laws as counterfactually robust as we had hoped.

Loewer recognizes this, and the view he arrives at is closer to the permissive imperialist view: “The special science laws that hold at t are the macro regularities that are associated with high conditional probabilities given the macro state at t” (Loewer, 2008: p. 21). “As the universe evolves... the probability distribution conditional on the macro state will also evolve.” We can illustrate this with our FISA as follows: she starts out with credence 1 in the laws of physics, and in the low-entropy macrocondition. Her conditional credences are uniform with respect to the those microstates that realize the low-entropy macrocondition. As the universe evolves, our FISA conditionalizes on macroscopic information–that is, information about the positions of middle-sized dry goods, their temperatures and densities, locations and velocities. At any time, having conditionalized on all of the universe’s macroinformation, those generalizations with high probability are the special scientific laws at that time.

Here we have a permissive imperialist view with a well-defined notion of robustness: the robust initial conditions are those which are encoded in the world's macrostate. But we can see immediately that this too is problematic: first, not all true macroscopic generalizations are laws; but all true macroscopic generalizations will get probability 1 on the scheme advocated by Loewer. Second, some true macroscopic generalizations will be laws despite *not* having high probability conditional on macroscopic information. TAke our generalizations of population genetics. Presumably these are true because of some facts about the structure of the chemicals which convey our genes. But these chemicals are *not* macroscopic; they are microscopic. So they will not be conditionalized on by our FISA, and the generalization will not be a law.

Perhaps there is a way of tweaking the Albert/Loewer view to account for this; my worry that there is no independently specifiable set of facts such that conditionalizing the uniform distribution over microstates on these facts will yield a high probability to all and only special scientific laws.

And this generalizes: for a permissive imperialist view to work, there must be some non *ad hoc* way of specifying which initial conditions are 'robust' enough to ground higher-level laws. Without such a specification, the imperialist has no way to distinguish laws from non-laws at the higher level. And without a way of distinguishing the laws from non-laws, we will not have the beginning of an explanation of COUNTERFACTUAL ROBUSTNESS and METHODOLOGICAL INDEPENDENCE. In order to explain why the special scientific laws are supported by induction and support counterfactuals, we must first distinguish between them and the non-laws, which are *not* supported by induction or counterfactually robust. The permissive imperialist cannot do this.

The Anarchist

In this section I'll briefly review some specific anarchic views. I'll then present some challenges to anarchism. First, anarchist views are not well-positioned to explain the fact that lower-level laws seem to trump higher-level laws: exceptions to the laws of economics can be explained by physics, but not *vice verse*. Second, anarchic views face a dilemma: either they accept that the *facts* (but not laws) of the special sciences depend on physics or they do not. If the former, they must provide a role for the special scientific laws in a world in which physics seems to do everything. If the latter, they must explain how independently operating laws produce a unified world. They have thus far been successful at neither.

The anarchist holds that the laws of the special sciences are laws for the same reason that the fundamental laws are. What makes the special science laws lawful? This question will be answered differently by different anarchists–Humean anarchists, like Craig Callender and Jonathan Cohen (2009, 2010), claim that they provide the best systematization of facts in the language of their science (though, for Callender and Cohen, the choice of language is arbitrary or pragmatic). Anti-Humean anarchists, like Nancy Cartwright (1997), hold that the laws of the special sciences, like the laws of physics, encode dispositions or capacities which manifest in the controlled environments that that science studies. There are, according to Cartwright, no principles coordinating the laws outside of these controlled environments.

While Callender and Cohen and Cartwright agree that the laws *and* facts of the special sciences and physics depend on one another symmetrically if at all, this is not a requirement of anarchism. I will call this breed of anarchism 'radical anarchism'.

According to the radical anarchist, our FISA will have a number of distinct, possibly incomplete credal functions available to her. Each of these will be defined over a different set of propositions: $F_{\text{biology}}(A|B)$, $F_{\text{physics}}(C|D)$. . .., where A and B are couched in the language of biology, and C and D are couched in the language of physics.

According to Callender and Cohen, A and B, C and D are different propositions because they come from different ways of partitioning the space of worlds; there may be some overlap between, say, A and C, and there may even be a translation between the AB partition and the CD partition, but the probability functions are distinct and defined over different propositions. Which credal function FISA uses depends, according to Callender and Cohen, on which is easiest for FISA to apply to the situation at hand. Which evidence propositions are most easily verified in this situation? Which conditional probabilities easiest to calculate?

Similarly for Cartwright, FISA will avail herself to a variety of disjoint credal functions, but instead of each being complete over a partition of the space of worlds, they will each be incomplete and only defined within certain controlled situations. So in situations in which $F_{\text{physics}}(A|B)$ is defined, $F_{\text{biology}}(A|B)$ is not. The situations in which physics yeilds a conditional probability are those with x-rays and scanning-tunnelling microscopes; the situations in which biology yeilds conditional probabilities are those in which groups of animals interact. Which credal function FISA uses will depend on the situation in which she finds–or creates for–herself.

It is compatible with anarchism that the *facts* at the special scientific level depend asymmetrically on the *facts* at the fundamental level; but anarchists deny that the *laws* so depend. Views of this latter sort–according to which the laws are in some way emergent, despite the dependence of the facts at the higher level on the facts of fundamental physics, are held by Fodor (1974), Lange (2009), and Armstrong (1983). According to these philosophers, the independence of the higher-level laws arises because the laws of the special sciences describe patterns which are visible only at the coarse-grained higher level, or are not the result of the laws of physics alone, or are the result of the laws of physics together with any suitably special initial conditions, or are backed by modal facts (necessitation relations or irreducible counterfacts) which are independent of both the lower-level modal facts and the higher-level categorical facts. Because this version of anarchism allows some dependence between facts at different scientific levels, we will call it 'moderate anarchism.'

Both varieties of anarchism score well in accounting for the METHODOLOGICAL INDEPENDENCE and, at first brush, the COUNTERFACTUAL ROBUSTNESS of the generalizations of the special sciences. The counterfactual robustness of special scientific generalizations is explained the same way as the lawhood of fundamental generalizations: either with *sui generis* modality or in terms of unificatory power.

Similarly,the methodological independence of the special sciences is explained easily by the metaphysical independence of the laws. Special scientists are able to perform inductions in the same way physicists are because their laws are metaphysically the same as those of physics.

Radical anarchism does poorly in accounting both for the MUTUAL CONSTRAINT and the ASYMMETRY of the special sciences and physics. On Cartwright's view, any two sciences don't attempt to describe the same world; rather, the make predictions about distinct controlled situations. No rules govern how they interact with one another, but plausibly the capacities of any science can overturn those of any other. So it's surprising that scientists seek information from one another, and that contradictory predictions are taken to indicate that one or another science's laws must be altered.

Radicals realize this; both Callender and Cohen and Cartwright argue that neither of these hold[36]. Unfortunately I do not have space to address their arguments here; so we will give them a demerit for failing to account for these relations, but note that this consequence of their view is not one these folks take to be a negative.

Moderate anarchism does better in explaining MUTUAL CONSTRAINT and ASYMMETRY of the coordination problem. According to these views, constraint and asymmetric dependence arise from the metaphysical dependence of the *facts* of the special sciences on the facts of

---

[36]Callender and Cohen reject asymmetry, but accept mutual constraint. On their view, each science forms a deductive system in an independent vocabulary. Because the vocabularies describe the same world, they must agree on the categorical facts of the world. Consequently no generalization at any level can imply that another generalization is (actually) false. However, nothing in their view guarantees that the laws will agree on what happens in counterfactual situations: a systematization could rule that, for some merely possible event A, if A were to happen, then B would, while another could rule that if C were to happen, then D would, where A metaphysically entails C but B and D are mutually contradictory; if A does not occur Callender and Cohen can't guarantee that this would not be the case. Similarly, they cannot guarantee that the chances assigned by various laws will yield compatible constraints on credence.

fundamental physics. We were understandably mistaken in our belief that these constraints held at the level of laws.

This view cannot be correct. For the laws of the special sciences have exceptions, and these exceptions can be explained by the laws of lower-level sciences. In fact, in many (though not all) cases, the exceptions to a special scientific law can only be specified by appeal to a lower-level science. Whether or not the laws of the special sciences have *built in* ceteris parabis conditions , frequently specifying situations in which they do not hold requires us to take on board concepts *which are not a part of the special science in question.* The predictions of economics can be trusted *provided* an asteroid does not strike the market.

A more subtle example from biology involves the Hardy-Weinberg law, which says that, in absence of evolutionary forces, the alleles and genotypes in a population will remain constant. This law has exceptions, but explaining *for which* species the Hardy-Weinberg law holds can only be done by discussing properties of DNA. Similarly, explaining which–highly unlikely–scenarios are entropy-increasing and so violate thermodynamics' second law can only be done only by citing the momenta of the particles underlying the system.

But meteors impacts are not describable in the conceptual scheme of economics (we have astrophysics for that), describing DNA proteins requires chemical, and not merely biological, concepts, and discussing the (non-aggregate) features of the particles which make up a gas is outside of the conceptual sphere of classic thermodynamics.

This observation allows us recognize a problem for anarchism's explanation of COUNTERFACTUAL ROBUSTNESS. For though the anarchic view may be able to explain the force of the special scientific laws, it is unable to explain why their exceptions are often outside of the conceptual scope of the science in which they feature. The anarchist has no explanation of the fact that lower-level laws trump higher-level laws and lower-level laws can explain the exceptions to higher-level laws, but not vice versa.

All versions of anarchism face the *conspiracy problem* (see Callender and Cohen 2010 for a discussion). If the laws of physics and the laws of the special sciences are independent, how is it that they conspire together to produce a unified world? That is, why is that the laws of physics somehow 'know' not to push elementary particles around in a way which violates the laws of the special sciences? And how do the special scientific laws, like those of psychology, fail to license violations of the laws of physics? The conspiracy problem is a challenge to the anarchist solution to MUTUAL CONSTRAINT; the anarchist claims that the sciences describe the same world; but if she is radical and holds that their laws are metaphysically independent, how do they combine to create a coherent world?

A distinct, challenge for both anarchist views–but especially the moderate anarchist–lies in explaining the COUNTERFACTUAL ROBUSTNESS of the special sciences. We gave anarchists a strong pass on this earlier: their explanation of special scientific lawhood is, presumably, the same as their account of fundamental scientific lawhood.

Together, these problems create a dilemma for anarchist views. For the moderate anarchist: if the fundamental laws govern the fundamental facts, and the fundamental facts explain the special scientific facts, what is left for the laws of the special science to do? For the radical anarchist: if the laws all independently determine the facts, how do they manage to produce a consistent world? The more radical an anarchist is are, the less she can explain MUTUAL CONSTRAINT. The more moderate she is, the less she can account for the COUNTERFACTUAL ROBUSTNESS of the special scientific laws.

The anarchist response is to claim that the special scientific laws *explain* in a way which is not reducible to the laws of physics. But note that this requires us to (a) take lawhood to be deeply tied to explanation, rather than governing, and (b) accept an overdetermination of explanation. While many philosophers, especially of the Humean strain, will not find either of these especially troubling, philosophers who take lawhood to be connected with governing,and who take explanation to be similarly tied to causation (rather than unification), will find this especially troubling.

The Democratic View

I've argued that a successful solution to the coordination problem cannot take the lawhood of the special sciences to be dependent on the laws of the physics. And I've further held that the laws of the each science cannot be made laws entirely be facts within the domain of that science. Both views leave at least one of our explananda–METHODOLOGICAL INDEPENDENCE, COUNTERFACTUAL ROBUSTNESS, MUTUAL CONSTRAINT, and ASYMMETRY unaccounted for. How, then, can these desiderata be met?

In this section, I'll present a view according to which the sciences work together to generate a unified body of knowledge. The generalizations in any science are laws, not because of their explanatory capacity given the facts of that science, or because of their relation to more fundamental generalizations, but because of their contribution to the informativeness of the total set of scientific laws. The mutual constraint the laws exercise on one another is a result of the fact this informativeness is evaluated holistically: the laws of all sciences taken together contribute to the informativeness of our system. So they need produce an internally consistent and mutually reinforcing set of predictions. And the independence of the various sciences is also accounted for: each science contributes laws to the overall system independently.

The Democratic Best System

The view I'll defend here builds on the view of laws defended in ch. 1. Recall that the motivating story behind this view is this: the laws of nature are those generalization which will be most useful to us as agents embedded in the world. We are born ignorant, with a finite memory and a limited sensory apparatus. Despite these obstacles, we are able to learn a surprising amount from our environment; we use the laws to encode those things we can learn in one situation and bring to bear in another. In ch. 1 I identified five virtues a lawsystem must have to be useful to us in this way: breadth, strength, fit, modularity, and simplicity.

I'll now focus on strength and breadth. In §[ERA] I briefly developed a local notion of strength which allows us to capture the information a set of laws gives us about subsystems of the world. Recall that while we are interested in finding the most informative lawbook we can, we prefer that this information come in the form of widely applicable dynamic laws–laws which operate as functions from states of the world at one time to states of the world at another. Our laws need to be as informative as they can be, compatible with their being discovered and confirmed through repeatable experiments. The laws are those most informative generalizations which we can formulate by repeated observation. Our preference for dynamic laws is explained by their repeatability: dynamic generalizations, but not initial conditions, can be observed in action over and over. Similarly, we prefer simpler laws because they can be more quickly supported evidentially and because they provide more accurate predictions.

Remember, though, that we don't come to new situations with our hands empty. Rather, we are able to bring to the table some knowledge of about the system we're observing, manipulating, or explaining. We then look to the laws to provide us with a *lot more* information about the system. This can be represented by introducing a slightly more complicated way of understanding the breadth and strength of our lawsystem.

To see how this bears on the coordination problem, let's return to our FISA and consider her interests in formulating lawful generalizations. She is interested in discovering the most informative set of *conditional* probabilities $F(P|B)$,where P is a prediction and B is a set of boundary conditions. Her lawbook can be strong in two ways: first, it can be strong by being accurate: the conditional probabilities can be such that $F(P|B) \approx 1$ for situations in which P and B, and $F(P|B) \approx 0$ when B and ~P. But her lawbook can also be strong by being more applicable: that is, it can give her predictions for a wider range situations, represented by the boundary conditions B. Call the first variety of strength *accuracy*, and call the second *comprehensiveness*

Accuracy and comprehensiveness trade off against one another: a lawbook can gain comprehensiveness by applying to situations with less uniform phenomena, although by doing so it will be unable to provide as accurate predictions of their behavior. Maximizing the combination of these virtues is hindered by the fact that the laws need to be *repeatable*:

they must be formulated in such a way that multiple distinct situations have, according to the laws, the same boundary conditions, and the laws must yield the same predictions in situations with the same boundary conditions. This is a requirement if the laws are to be discovered and evidentially supported by induction. The laws are generalizations which we can learn in one context and apply to another.

Recall from §[ERA] that a system can be broad in two different ways: by applying to many QISs, or by applying to a variety of quite different QISs. This difference is now playing an important role. The first way in which a lawsystem can be broad, by having many QISs, is here called 'repeatability'. The second way, by having quite different QISs, is here called 'comprehensiveness.' We can have more accurate probabilities that are tailored to each experimental situation, but they will not be repeatable; we can have a probability functions which is highly accurate but only by excluding some situations, but it will not be comprehensive; and we can have a repeatable, broad probability function that moves further away from 1 for some true predictions and further away from zero for some false ones.

Fundamental physics is extremely accurate. But it is not maximally broad. For any maximally fine-grained propositions *B* and *P,* a deterministic physics will give zero to P if and only if P is false and assign one to P if P is true. But physics is silent about less coarse-grained propositions: let B be the proposition that the temperature of a gas is *T,* its volume is *V*, and its pressure is *P.*

Our agent's information about the boundary conditions of a system need not be maximally fined grained. But if she conditionalizes on the proposition that some system's pressure and volume increase, what does fundamental physics say about the gas's temperature? Unfortunately, nothing. For even adding *a posteriori* identities relating physical properties to thermodynamic properties, we still will not arrive at a prediction for the temperature of the gas: there are physical states compatible with the boundary conditions which are temperature increasing, and physical states compatible with the boundary conditions which are not. So while a set of laws in terms of maximally fine-grained propositions may be accurate, it will not be broad.

To increase the comprehensiveness of the laws, we may add laws which take us from course-grained states–like temperature and pressure–to other course grained states. This is project of thermodynamics. Or we can add a probability function over the fine-grained states which is invariant under the fine-grained dynamics. This is the project of statistical mechanics. In either case, our predictions will diverge from perfect accuracy, so we will lose some accuracy in the overall system. But we will gain comprehensiveness.

I claim that each scientific discipline increases the breadth of the overall lawbook. By adding more laws at a higher level, we increase the comprehensiveness of the overall system with some moderate sacrifice to its accuracy. The view here builds on the work of Handfield and Wilson (2013)[37].

Let's see how our FISA will behave on this way of understanding the laws. She will begin with a set of fundamental laws; she'll work out the consequences of these laws, and

[37]Handfield and Wilson deliver an apparatus for combining distinct objective probability functions at various levels of grain without generating the sort of contradictions described in Meacham (2013), but they do not offer a metaphysical view of probability to motivate their hierarchy. The view described below provides a motivation for the sort of heirarchical view described by Handfield and Wilson and extends the account to deterministic laws.

generate a probability function $F_0(P|B)$. This probability function will be incomplete; it will only be defined for maximally fine-grained propositions *P* and *B*. So FISA will see if there is a set of more coarse-grained variables in which she can formulate fairly accurate and repeatable laws. She'll add these to her lawbook, and work out the consequences, arriving at an extended credal function $F_1(P|B)$. But this credal function still won't be defined over propositions at all levels of grain, either because these coarse-grained laws don't imply *more* coarse-grained laws or because these implications are cognitively intractable for FISA (recall that FISA is not logically omniscient; she, like us, finds some inferences to complex to complete). So she will find another set of repeatable yet accurate generalizations at a coarser level of grain and *these* to her lawbook, generating $F_2(P|B)$. When does this stop? Whenever FISA either FISA's credal function is defined over all propositions (unlikely) or she's unable to find laws that have an acceptable degree of both repeatability and accuracy.

Let us make this more precise. We require our lawbook to be formulated in terms of a series of variables. Setting all of these variables determines a *state* of the world. The variables thereby partition the space of nomically possible worlds, with each cell of the partition corresponding to a unique state of the world[38]. This requirement is not motivated by considerations of fundamentality (as is Lewis' naturalness constraint). Rather, to be repeatable and comprehensive our lawbook must identify some situations as identical with

[38]Thus far, nothing prevents the two worlds from differing without differing with respect to the quantitative properties the laws concern. We should take this to be a benign consequence: for if physics is complete, each cell in the partition induced by the variables of physics contains exactly one world. But if it is not, some worlds differ without differing physically, and so have the same physical state despite being distinct. Whether or not physics–or any science–is complete in this sense should be expressible by our theory of laws but not determined by it.

respect to the quantities about which it yields predictions, and its predictions must be functions of those quantities.

But our information about boundary conditions can vary in its degree of precision; less precise information is a coarse-graining of more precise information. The fundamental physical laws, together with boundary conditions specifying the heat and volume of a gas, yields no predictions about the gas's future state. This gives us reason to include both coarse and fine-grained variables in complete lawbook. Given a set of fine-grained variables V, we can expand to include a more coarse grained set by adding variables V' which are such that, for each state S identified by setting the variables of V, there is some state S' defined by the variables of V' such that S ⊢S' but not vice versa. If this is the case, then each cell of the partition induced by the variables of V' will be a disjunction of cells induced by the variables of V. We can call the union of these two variable sets V+.

These variables will either represent the fundamental quantitative properties of the world or be coarse-grainings of variables that represent the fundamental quantitative structure of the world[39]. We can now evaluate law-systems which include information at different levels of grain.

---

[39]As it stands, this requirement on our variable space looks similar to the naturalness constraint of Lewis (1983). But it is just suspicion about this constraint that let Callender and Cohen to reject the traditional best system! Appealing to it here is suspicious at the least. But we can drop the requirement that the variables be metaphysically fundamental; instead, we can hold that the laws identify some set of variables as fundamental, and that all laws are coarse-grainings of these (nomically, but perhaps not metaphysically) fundamental variables. We can then require that our lawbook contain a set of macroscopic variables in which we are particularly interested as a coarse-graining of its fundamental variables, and evaluate its informativeness in tandem with its terms using a method similar to that of Loewer's (2007) Package Deal Account.

We will evaluate the informativeness of a lawbook by evaluating the probability function it generates. But to do so, we need a recipe for generating a probability function from the laws. We will do so as follows: if $A \rightarrow B$ is *derivable* from our lawbook, then $P(A|B) = 1$; if $P(A) = f(x)$ is derivable from our lawbook, where $f(x)$ is some function of our variables, then $P(A|B) = f(B)$, where B is a proposition giving the values of the variables in *x.*

The notion of derivability here is importantly weaker than implication; for austere lawbooks will imply many facts which are cognitively inaccessible to agents like us because of the computational complexity involved in deriving them. In such a case, we can and ought to add higher-level laws by hand, even if these laws decrease the overall accuracy of the lawbook. Rather, it is something closer to *cognitive accessibility:* a generalization is derivable in the sense here specified if sufficiently idealized scientists could derive it. This adds a parameter to our theory of laws: if our scientists are *very* idealized, then they will be able to generate a more informative probability function from fewer laws; if they are more like us, they will need manually increase the informativeness of the lawbook by adding more laws. Call this parameter *accessibility,* where a lawbook is more accessible when its laws are derivable by less ideal scientists[40]*.* Because not all of the implications of our laws

[40]Rather than being a disadvantage of the view, the addition of this parameter allows us to generate a hierarchy of lawlikeness. For some laws will only feature in the most accessible lawbooks; these laws are more approximate than those which feature in the least accessible, most austere lawbooks. Take, as an example, the laws of classical mechanics. These laws ar only approximately true, but they somehow manage to support counterfactuals, appear in explanations, and underwrite predictions. Nonetheless they are in some way less deserving of the name ’laws’ than the laws of quantum mechanics. On my view, this is because the laws of classical mechanics are unneccessary for agents with access to the laws of relativistic quantum field theory and unlimited cognitive capabilities, but extremely useful and informative for agents more like us, who find the equations of quantum feild theory impossible to solve exactly except in very simple situations.

are accessible, our probability function will be incomplete: it will not be defined for boundary conditions which, when taken as inputs to the equations of our laws, yield equations too complex to be solved by our less-than-ideal scientists.

The accuracy of the lawbook is evaluated as follows: each lawbook is given an *accuracy score* using some scoring rule function[41]. The probability function the lawbook generates is a conditional one; to evaluate it accuracy, we look at situation which have the boundary conditions described by the laws. We update the conditional probabilities on those boundary conditions, and see what probability the laws assign to the actual outcome of the situation. The closer the probability assigned by the laws is to the actual outcome, the higher the laws score on accuracy.

The comprehensiveness of the laws is determined as follows: given the lawbook's accessibility, over how many *actually instantiated* propositions is it defined? Recall that it need not be defined for *all* actual situations, nor need it be defined at all levels of grain. The more situations and levels of grain over which it is defined, the more comprehensive it is.

This is the democratic view: each science represents a distinct level of grain, at which we must balance accuracy and repeatability to formulate laws. But the justification for adding new sciences is to improve the overall score of FISA's credal state in terms of accuracy, repeatability, and comprehensiveness. How, then, do we satisfy our four requirements?

---

[41]For more detail on scoring rules, see Joyce (1998), (2009), and Leitgeb and Pettigrew (2010). Although some features of scoring rule functions are agreed upon–and these are features of quadratic scoring functions–there is single agreed-upon function.

METHODOLOGICAL INDEPENDENCE: The laws of each science are added to the lawbook because they individually increase the informativeness of the lawbook. Determining which generalizations will fill this role at some level of grain is the job of each special science.

COUNTERFACTUAL ROBUSTNESS: The laws of each science are laws for the same reason: they increase the comprehensiveness of our system of laws without weakening its accuracy or repeatability. They support counterfactuals for the same reason the fundamental laws do. Of course, for a Humean, this story is complicated; the short version says that the laws are counterfactually robust because they ground coutnerfactuals by fiat; the longer version justifies this stipulation by the pragmatic utility of holding these particularly informative and supported generalizations fixed while evaluating counterfactuals.

MUTUAL CONSTRAINT: Because the accuracy, comprehensiveness, and repeatability of a law system is evaluated holistically, we can expect the laws not to contradict one another–if they did so, the accuracy of the lawbook would be obviously compromised, and we should expect the various sciences to inform one another. Discovering connections between sciences allows us to insure the mutual consistency of our overall belief structure.

ASYMMETRY: The facts at each level is a coarse-graining of some lower level. If B is a coarse graining of A then setting the value of A determines the value of B (but not vice versa). So more fined-grained information screens out more coarse grained information, and the facts of the higher level science are implied by the facts at the lower level sciences[42].

[42]It is just this asymmetry that requires us to add special scientific laws to our system: though the fine-grained information settles the coarse-grained states, coarse grained boundary conditions tell us nearly nothing about their fine-grained realizers. So we need add higher-level laws to make predictions given coarse-grained information.

I conclude that the democratic view neatly explains all four features of the coordination problem: methodological independence, lawhood, mutual constraint, and asymmetry. Its explanation of METHODOLOGICAL INDEPENDENCE and COUNTERFACTUAL ROBUSTNESS are reminiscent of the radical anarchist; its explanation of MUTUAL CONSTRAINT and ASYMMETRY are close to those of the imperialist and the moderate anarchist, respecively. In this way it poaches the best features of each of the views I've discussed.

Further Advantages of Democracy

It's worth noting here that even on the most austere inaccesible lawbook, the laws of fundamental physics will not be wholly comprehensive. For while they will be defined over all fine-grained propositions, they will not have any defined probabilities conditional on coarse-grained information. For a coarse-grained proposition is a collection of infinitely many finely delineated microphysical states; there are infinitely many arrangments of fundamental particles corresponding to the proposition that *the heat of this gas is forty Kelvin*, for example, and nothing about that proposition gives us reason to take any of its microphysical underlyers to be more likely than any others.

Albert (2000) and Loewer (2009) argue on the basis of this sort of consideration that the lawbook must contain PROB, a law specifying a probability distribution over initial states. Such a distribution will yeild conditional probabilities conditional on any macroscopic proposition compatible with microphysics. I've argued previously that Loewer and Albert do not go far enough because they cannot account for the lawfulness of special science generalizations; the view I defend here justifies the inclusion of the laws of the special sciences by appeal to the cognitive intractability of deriving conditional probabilities from

PROB for most special science generalizations. Interestingly, though, on my view a creature twofold: first, Lewis' notion of strength doesn't allow with unlimited cognitive capacities would be interested in including PROB in her lawbook, but no other special science laws. So on my view PROB has a special status.

This view of laws neatly accounts for two other features of special scientific laws which have been recognized by various authors (Mitchell (2000), Woodward (2003, 2013)): first, the laws of special sciences have exceptions, but these exceptions cannot be captured in ceteris parabis clauses using the concepts of the special science (Woodward (2003), Cartwright (1997)). On the view sketched, special scientific generalizations are lawful if and only if they feature in a system which acceptably balances accuracy, comprehensiveness, and repeatability. Laws which have exceptions can lack perfect accuracy but, by being repeatable and extending the comprehensiveness of the system, be worthwhile additions to the lawbook. Their inclusion does not require their exceptionlessness, nor does it require that there be formulable or nonredundant ceteris parabis conditions limiting their scope.

Secondly, the worthiness of special science vocabulary is not dependent on its definability in fundamental terms. On Lewis' view, whether a term is eligible for use in a special science depends on its degree of naturalness; degree of naturalness depends, for Lewis, solely on the length of its definition in perfectly natural terms. This means, among other things, that the predicate 'electrino', which we stipulate to refer to electrons created before 2015 and neutrinos created after 2015, is more eligible to feature in a special scientific law than is the term 'mammal,' which presumably has an extremely complex and disjunctive definition in perfectly natural terms. ’Electrino’ is not more natural than ’mammal’, and independently of our view of special science vocabulary we should recognize that mammals are more similar to one another than electrons are to neutrinos, whenever they are created. On the view here offered, the eligibility of a term instead depends on whether the comprehensiveness of a set of laws can be sufficiently increased by adding laws in those terms to our complete lawbook[43].

While it is a requirement of the view that the higher-level terms force a partition of worlds which is a coarse-graining of those offered by the lower-level terms, this minimal constraint does not make the relative eligibility of coarse-grainings dependent on anything other than the informativeness of the laws so phrased, as measured by accuracy, comprehensiveness, and repeatability.

Thirdly, laws come with various degrees of lawfulness. Some laws are less modally robust: they hold in fewer situations and they are less stable than others. There is a continuum of laws, starting with the laws of physics, which are exceptionless and maximally modally robust, moving through the central principles of special sciences, like the principle of natural selection or the thermal relaxation time of a certain sort of liquid, and culminating in mere accidental generalizations. The view sketched, unlike the dispositional account of laws, has the capacity to account for this. For there are more than one way to weight the three virtues this view rests on: if perfect (or near-perfect) accuracy is given maximal weight, then only the laws of physics are included in the lawbook. By varying our

[43]For a more in-depth discussion of the difficulties involved in tying the Lewisian notion of naturalness to our account of laws, see Loewer (2007) and Eddon and Meacham (2013); for a discussion of this problem focusing on special scientific laws, see Callender and Cohen (2009).

permissiveness for accuracy, we will vary the generalizations which are permitted in the lawbook. Those which count as laws on more accurate rankings occupy a more privileged place on this continuum than those which do not.

This hierarchy can be tied to the counterfactual robustness of the laws. For the view sketched is Humean, according to which counterfactuals are made true by the laws. Plausibly[44], the strength of counterfactual support varies with the accuracy of the laws[45]. So counterfactuals which are made true by the laws of physics, our most accurate set, override those made true by biology. And within a science, the counterfactuals made true by more accurate laws trump those made true by less accurate laws–so the counterfactuals made true by quantum mechanics trump those made true by classical mechanics.

I've claimed that a particular form of the Best Systems Account of lawhood can explain the relevant features of the relationship between laws in various scientific disciplines. Can a more metaphysically robust view do this? I am doubtful. For a key feature of the view is taking the informativeness of the lawbook, measured in a particular way, to be partially constitutive of lawhood. Anti-humeans reject the claim that the laws are, by their nature,

---

[44]or by stipulation!

[45]Woodward (2003), 6.12, argues against Sandra Mitchell's (2000) notion of *stability* and Brian Skyrms (1995) notion of *resiliency* on the basis that these nonmodal notions cannot capture what we are really interested in in discovering laws, *vis*, their counterfactual stability (this point is also made by Lange (2009)). On the view offered, as in other Humean views, counterfactual stability is grounded in occurant facts, in this case, a sort of stability across situation in a similar vein to that described by Mitchell and Skyrms. So it would be a mistake to criticize this view for missing the counterfacts–they are true because of the occurant facts described. Of course, the proof is in the pudding: does the sort of stability here described generate the *right* counterfactuals? I hold that it does.

informative. So no way of measuring the informativeness of the lawbook will suffice to make some higher-level generalization a law.

Nonetheless proponents of metaphysically robust views who hold that *only* the fundamental laws are backed by modally robust fundamental facts can appeal to the view I've defended to distinguish between accidents and laws at a higher level. Many philosophers who doubt that a fully Humean story can be told about the fundamental structure of the world are subject to the criticisms laid at the feet of the imperialist in [The-Imperialist]. Although the view that results will have a different explanation of the counterfactual robustness of the special science laws from that of the laws of physics, they will inherit the other advantages of the democratic view.

Conclusion

Extant views describing the relationship between distinct scientific disciplines leave key features of this relationship unexplained. This failure manifests itself in philosophical views about the lawhood of special scientific laws; these views, no matter their metaphysical commitments, fail either to account for the autonomy of the special sciences or for the mutual dependence of scientific disciplines. A Humean view, which takes the informativeness of the laws to be partially constitutive of their lawhood, measures informativeness by the accuracy of predictions made by the laws on the basis of repeatable boundary conditions, and evaluates the informativeness of all sciences together, in uniquely able to capture these features of the relationship between laws. I then point out additional advantages of the view: it accounts for the degrees of lawfulness of special scientific laws, the fact that special scientific laws have exceptions, and the fact that explaining these exceptions often requires concepts that are not a part of the science in which the law is formulated.

The view outlined in this chapter is a development of the account of laws defended in ch. 1; it builds on that view by tying our account of the informativeness of the laws more directly to our interests and uses for them. The accounts of breadth and strength in ch. 1 dealt with unstructured propositions and differentiated subsystems of the world via their intrinsic properties. The accounts offered here tie breadth to the sort of information about boundary conditions embedded agents have access to, and tied strength to the amount of information agents in a position to use some information could extract from the laws. I view this as a development of the ideas presented in ch. 1. owever, the central presented in this paper, that the laws of the special sciences contribute to our lawsystem by allowing us to evaluate predictions and counterfactuals when we have too little information to employ the fundamental laws, can be separated from the view defended in ch. 1, and either added to a metaphysically robust account of fundamental lawhood, or seen as providing a role for the special sciences distinct from that defended for fundamental sciences in ch. 1.

Making Fit Fit

Introduction

Thus far I have been reluctant to discuss the humean account of objective probability. The humean account of objective chance relies on a notion of *fit*, which I have been content to leave vaguely understood. But the view of laws which has motivated the previous two chapters also provides us with some insight on humean accounts of chance.

Objective chances have two roles. First, they are part and parcel with the laws of nature. With the laws they play a role in determining–probabilistically–how the world unfolds. But they also have a normative role. Chances ought to impact our partial beliefs, or credences. Our credences should match the chances.

Humeans about objective chance have followed Lewis (1980) in understanding objective chances through their relationship with the laws of nature and in taking principles linking chance to credence to be central to our understanding of objective chance. The humean theory of laws advocated here is well-suited to connect the two roles of chance: laws, according to a Humean, are true generalizations that sit between induction and prediction. Objective chances, thinks the Humean, are just like laws, except for them informativeness is measured by their impact on credences rather than on full belief.

Lewis (1980, 1994) calls the measure of informativeness for chancy laws 'fit'. For Lewis, laws *fit* a world in proportion to how likely they make that world. This way of understanding the informativeness of probabilisitic laws is flawed: it offers too mediated a connection between the chances and our credences. In this chapter I employ notions from epistemic utility theory to advance a better measure of the informativeness of probabilistic laws. The view developed has two primary advantages over the traditional Lewisian view: first, it provides a more direct explanation of the fact that chances are probabilities; second, it provides a simpler explanation of the normative force of chances on credences. On the view advanced here, the chances are better fit to play their normative role than on the Lewisian proposal.

Recall that the view put forward in this dissertation is inspired by the following guiding principal: Humean laws are those generalizations which sit between induction (from experience) and deduction (of prediction, retrodictions, and conditionals). Traditional Humean views–those of Hume, Mill, Ramsey, and Lewis–have focused on the deductive aspect of laws[46]. The laws, on this view, are the best way of organizing all facts, whether or not those facts or the laws are accessible to agents in the world or scientists.

The criticisms of and refinements offered in this dissertation aim to bring the laws down to us. The chief advantage of the ideal Humean view is that it shows why embedded agents have a use for laws, causation, and counterfactuals. The role laws have to play here is in extending our knowledge from the observed to the unobserved. To do this, they must be epistemically accessible to agents embedded in the world. Such a system is better, not just if it contains more information, but also if it is more conducive to empirical discovery. This requires us to look for generalizations which can be easily discovered through observation as well as easily extended to prediction.

[46]So Ramsey says "Even if we knew everything, we should still want to organize our knowledge in a deductive system [...] what we do know we tend to organize in a deductive system and call its axioms laws" (Ramsey, 1927).

The chapter is structured as follows. In §[BSA and PP] I provide a critical overview of the Best System Account theory of laws, with a focus on the BSA's virtue of fit and its relationship to strength and simplicity. I'll then discuss the form of chance-credence linking principles. The chance-credence link I will focus on for the majority of the paper is Ned Hall's New Principle (Hall (1994)). Taken together, these provide an opinionated introduction to the two aspects of chance. In §[Fit and Accuracy] I develop an accuracy-based account of informativeness for probabilistic laws. In §[Chance and Utility] I show how the metaphysical picture of chance defended here fits with one accuracy-based argument for the New Principal. I'll conclude by considering objections to my view.

The Best System and the Principal Principle

Objective chances have two faces. The first looks to causation, natural law, and scientific explanation. The chances make the world the way it is; they help explain why some (likely) events occur and others (unlikely ones) do not.

For example, an atom of carbon-14 has a .5 chance of decaying within 5,730 years. This fact is a *non-accidental* fact about carbon-14, it explains why half of a particular sample of carbon-14 has decayed given that that sample is 5,730 years old, and it grounds the causal production of a block of nitrogen-14.

The role of chances in explanation is similar to that of laws. The trajectory of the cue ball and the eight ball at t_0 explains their trajectories at t_1 *in virtue of* the law of conservation of momentum. The fact that the cards in the deck are half red explains that half of my draws are red *in virtue of* the fact that I was equally likely to draw each card in the deck[47].

Because of the similarities between law and chance in underwriting explanations and causal facts, we should expect our accounts of law and chance to be closely connected. In §[subBSA] I'll outline an account of chance and law that connects them: the regularity theory of law, which fits neatly with modified frequentist accounts of chance.

The second face of chance looks to our partial beliefs. If the chance of an atom's decay in the next twenty minutes is .5, I should be just as confident that it will decay as I am that it will fail to decay. But if I know that the chance of drawing a blue ball from an urn is .25, I would be irrational to take a draw of blue to be more likely than not. Because chances play a normative orle in guiding belief, we should expect them to be connected to traditional epistemic values, like evidence and truth-conduciveness

Precisely characterizing this connection is not trivial. In §[chance credence link] I'll outline some proposed principles linking chance to credence, and show how they fit with this account of laws. Later, in §[Fit and Accuracy] I'll argue for a view that connects the chances to truth and evidence.

Fit in the Best Systems Account

The regularity theory of law holds that laws are true generalizations; the frequentist account of chance holds that chances are actual frequencies. A naïve regularity account

[47]This explanation need not go by way of a principle of indifference; rather, the fact that I'm equally likely to draw each card is a result of my shuffling, which is a chancy dynamic process.

holds that all true generalizations are laws; a naïve frequentist account holds that the chance of (say) heads coming up in a coin flip is equal to the exact proportion of coin flips which result in heads.

Both naïve views are clearly false. Not all true generalizations are laws: some, like the fact that all of the eggs in my refrigerator are brown, are merely accidental. Laws support counterfactuals, but this generalization about my refrigerator does not: were I to buy white eggs, they would not become brown when I put them in the fridge. Similarly, not all chances precisely match their frequencies. For suppose the actual number of coin flips were odd; then it would be impossible for exactly half of the flips to be heads. But it is ridiculous to suppose that a fair coin could not be flipped an odd number of times. Similarly, the quantum mechanical probability of some events is an irrational number, but as actual relative frequencies are ratios of integers they cannot be irrational.

Refined frequency and regularity accounts circumvent these counterexamples by restricting the generalizations which are laws and loosening the link between chances and frequencies. The most developed such account is David Lewis' (1980, 1994) Best System Account (BSA) of laws and chances. Lewis's account holds that the laws and chances are those generalizations and chance-statements which form deductive system which maximizes three virtues: *strength*, *simplicity*, and *fit*.

We'll call a set of generalizations and chance statements a '*lawbook*'. Each lawbook receives a score for each virtue. Lawbooks get higher marks in strength for implying many true propositions (or, equivalently, ruling out many worlds). Lawbooks get higher marks in simplicity for being syntactically shorter: by having fewer and shorter sentences[48]. Lawbooks get high marks in fit for having chances which give a high probability to the actual world. The points in each category are tallied; the lawbook with the most overall points comprises the scientific laws of the world.

The BSA escapes the counterexamples to the naïve regularity and frequentist accounts by evaluating lawbooks holistically and invoking a simplicity consideration. Because lawbooks get higher marks for simplicity when they have fewer sentences, not every true generalization will be a law–only those whose contribution to the inferential power of the lawbook outweighs their cost in syntactic length. Similarly, Lewis suggests that a lawbook may contain probabilities that diverge from the actual frequencies for simplicity considerations. For example, a lawbook may assign 0.5 to a coinflip coming up heads despite the actual frequency of heads being 0.4999999.

A brief advertisement of this approach: if we take the laws to be the theorems of the best system, defined in Lewis' sense, we can give a boilerplate justification of their counterfactual robustness. The laws underwrite counterfactuals because they are organizationally central to our belief set: the laws are amongst a small set of beliefs (because of their simplicity) from which many other propositions in our belief set follow

[48]Understanding simplicity in terms of syntactic complexity makes our simplicity requirement language dependent. This means either we must hold that there is exactly one preferred language of evaluation or that generalizations are laws only relative to a choice of language. Lewis held the former; for an exploration of the latter view see Callendar and Cohen (2010, 2011). For a mixed view, on which the simplicity of a lawbook is evaluated both by its syntactic length and the complexity of the translation between the lawbook's preferred language and our language, see Loewer (2007).

(because of their strength). If we remove the laws from our belief set, this will generate a large change in the set.

When evaluating counterfactual or indicative conditionals, we employ the Ramsey test: we first add the antecedent of the conditional to our belief set and make as few changes as possible[49]. Because of the organizational centrality of the laws, removing them requires large changes in our belief set. So when evaluating conditionals, we hold the laws fixed.

Now for some brief criticism: the BSA unnecessary splits its measures of informativeness in two. Strength and fit both measure the inferential power of the laws: strength measures how much they tell us, whereas fit measures how accurate they are. There is very little difference between a law-system that gives a world a probability near zero and one that rules that world out completely; but measures that distinguish strength and fit see this as a great difference. While impossible worlds are assigned probability zero, not all probability zero worlds are impossible. For if there are infinitely many worlds (and there are) each world will have probability zero, but not all are impossible. These systems will be similar in their affect on our credences, but different in their affect on our full beliefs about nomic possibility. I take this difference to be small rather than large, as (I hold) the latter can

---

[49]Which beliefs we change depends in interesting ways on the type of conditional we're evaluating. When we evaluate ``if Oswald didn't kill Kennedy, someone else did" we change fewer beliefs than when we evaluate ``if Oswald hadn't killed Kennedy, someone else would have." In both cases, however, we hold the laws fixed. One task of a metaphysics of laws is to explain this constancy; the proponent of the BSA can do so by appealing to our conservatism in changing our beliefs, the epistemic centrality of laws, and our use of the Ramsey test in evaluating conditionals. Proponents of more metaphysically robust accounts of law must do so, apparently, by making brute stipulations about their metaphysical danglers.

typically be recovered from the former–for example, we can take the impossible worlds to be those which are in regions of probability zero.

Lewis' account of fit, like his account of strength, applies to worlds as a whole, rather than local events. Just as a system may be strong but allow inferences only from global states of the world to other global states, a system may fit well but provide little information about the likelihood of local events. Such systems are almost completely useless to embedded agents, who can neither discover nor apply them. This should worry us; our laws should be *ours*, both available and useful to agents like us. I'll make some suggestions for correcting this in §[infinity].

More troublingly, the BSA's simplicity constraint plays a vital role in distinguishing laws from accidents and in diverting the chances from the actual frequencies. But there are good reasons to doubt that Lewisian simplicity, measured by syntactic length, is fit to play this role. The complaint against the BSA's simplicity requirement here is threefold: firstly, simplicity is meant to explain why we don't take particular matters of fact, like the fact that there is exactly one liter of coffee in my thermos, to be part of the lawbook. But this is *not* the way in which scientists invoke simplicity. This moves the BSA a step further from its claim to scientific plausibility.

Recall from ch. 1 that scientists invoke simplicity considerations to decide between theories only when empirical considerations underdetermine theory choice, that is, only when they are comparing theories which are equally able to explain the phenomena under consideration. Thus we have no reason to believe that simplicity ever weighs *against* strength *or fit*. But to play the theoretical role it is required to, it must. This is the second complaint against the simplicity requirement.

Furthermore, it's not clear how simplicity should lead us away from the actual frequencies of the world. In what sense is .4999999 more complex than .5? Both are rational numbers. Perhaps we should count number of decimal places. But if this is correct we have a problem: quantum mechanics often assigns events irrational probabilities. These have infinite decimal expansions and so are by this measure maximally unsimple. But since irrational numbers are not ratios of integers, they cannot be the actual value of frequencies. Thus these single-case chances *both* depart from the frequencies *and* fail to maximize simplicity. So simplicity measured in this way cannot explain why chances diverge from frequencies.

I'll explore an alternative constraint in §[Fit and Accuracy]. The alternative constraint I offer is more closely tied to our epistemic access to the laws. Strength is counterbalanced in laws, not by a virtue measuring the ease by which they can be expressed, but instead by a virtue measured by how easily they can be discovered[50]. First, however, I'll provide an overview of the orthodox view of the second face of chance: its relationship to our partial beliefs.

[50]It's not obvious that this is best called 'simplicity', as that term seems tied to syntactic measures of complexity. But there is plenty of research on the curve-fitting problem and in Bayesian confirmation theory which argues that simpler theories are more quickly arrived at or provide more accurate predictions (see, for example **curve fitting shit** and **Rosenkrantz (1979) and Henderson (2013)**; for a Humean account which understands simplicity in this way see Hoefer (2007)).

The Chance-Credence Link

The connection between chance and credence is codified by chance-credence norms. Our partial beliefs should yield to the chances, but precisely how they ought to do so is a controversial matter. Three norms are currently live contenders: the Principal Principle (Lewis (1980)), the New Principle (Hall (1994), Thau (1994), Lewis (1994)), and the General Principal Principle (Ismael (2008)). I'll discuss each of these in turn, and then give reason to prefer the New Principle.

I'll follow Meacham (2010) in my presentation of these principles. Let $b(* \mid *)$ be a subjective initial credence function, let $ch_i(*)$ be some putative chance function and let T_i be the proposition that $ch_i(*)$ is the correct objective chance function[51]. Finally, let *E* be any admissible evidence–a notion which will be explained shortly. Then the Principal Principle is:

PP:

$$b(A|T_i\&E) = ch_i(A)$$

As a constraint on initial credences, the Principal Principle tells us to match our credences to the chances. That is, if we think that the correct chance theory assigns a chance of *x* to A, we should have a credence of *x* in A, whatever other (admissible) evidence we have. Lewis defines admissible evidence is any evidence whose impact on our credence in A comes entirely by way of the chance of A, but later provides sufficient conditions for admissibility:

[51]Meacham's presentation–and so mine–differ from Lewis' own presentation of PP: in Lewis' version, the chances are conditional not on the full chance theory, but merely the proposition that the objective chance of A is *x*. Nothing here hangs on this difference.

information about the laws and the chance theory of the world are admissible, for Lewis, as are any propositions about the past.

Unfortunately there is a problem combining the Principal Principle with the reductionist account of chances outlined in §[subBSA]. For recall that *fit* requires the chances to match the frequencies as much as simplicity will allow. Hence arbitrary mismatches between the chances and the frequencies are metaphysically impossible. As is customary I will call such mismatches *undermining futures* because they involve future states which would provide evidence against the chance function, and call chance functions which allow them *modest* chance functions. A modest chance function is so-called because it does not assign itself a chance of one–it is not certain that it is the right chance function.

As an illustration, suppose *T* is a chance theory that that assigns .5 to coinflips landing heads. If I know that the coin will be flipped exactly 400 times, what credence should I assign to *every* flip coming up heads? Because *T* will only be the chance theory if the frequency of heads is near 0.5, *T* implies that *all heads* is false. We should give no credence to the metaphysical impossibilities[52], so we should set $\mathrm{b}(\mathrm{allheads}|\mathrm{T}) = 0$. But if $\mathrm{ch}(\mathrm{heads}) = .5$, then $\mathrm{ch}(\mathrm{allheads}|400\mathrm{flips}) = 3.87 \times 10^{-121} > 0$. So by PP, $\mathrm{b}(\mathrm{allheads}|\mathrm{T}) = \mathrm{ch}(\mathrm{allheads}) > 0$. Hence the Principal Principle leads to inconsistent constraints on our credences: either we should have nonzero credence in the metaphysically impossible, or we should diverge our credences from the chances.

[52]It's open to the defender of chance reductionism to avoid the counterexample to the Principal Principle by rejecting this claim; surely one should not be dogmatically certain of one's metaphysical views. But we regard chance as a sort of epistemic expert. Whether or not we are certain of the metaphysical truths, chance, as an expert, ought to be. If it is, contradiction follows.

Nearly all reductionists about chance hold that the true chance function of the world *is* modest. Because the chances depend on the frequencies, the true chance function cannot assign a chance of 1 to its being true[53]. This leads reductionists about chance to reject PP in favor of one of two less naïve chance-credence principles: the New Principle (Hall (1994), Thau (1994), Lewis (1994)) and the General Principal Principle (GPP) (Ismael (2008)). First let's look at the New Principal:

NP: $b(A|T\&E) = ch(A|T\&E)$

NP tells us to set our credences to the chances conditional on all of our evidence, *including the fact that they are the chances*. In the special case of non-self-undermining chance functions, $ch(T) = 1$, $ch(A|T) = ch(A)$, and NP agrees with PP. But if a chance function assigns positive chance to propositions with which it is incompatible, $ch(T) < 1$, and so for undermining propositions U $ch(U|T) = 0 \neq x = ch(U)$. Thus $b(U|T) = ch(U|T) = 0$ without contradiction. In defense of NP, Hall (1994, 2004) notes that chance, via the Principal Principle, is an expert function–it is a credal function which we regard to be epistemically better than our own. Hall distinguishes between *database-experts* and *analyst-experts*. We defer to the former because we believe them to be better informed than we are; we defer to the latter because we regard them as better at *evaluating* evidence than we are. A database expert is someone you take to have more evidence than you do. An analyst expert is someone whose expectations you trust better than your own when you are both equally well informed.

[53]Although Jonathan Schaffer (2003) provides a recipe for generating immodest chance functions from modest ones.

Chance, according to Hall, is an analyst expert. The chances represent the best credences we can have based on our evidence. Consequently we should update them with our evidence before yielding to their advice. If our evidence includes the claim that they are the one true chance theory, we should let them know before matching our credence to theirs. Thus we should employ the New Principle rather than PP, which tells us to match the *un*conditional chances.

Finally, Ismael's General Principal Principle: GPP: $\mathrm{b}(\mathrm{A}) = \Sigma_{\mathrm{ch_i:chi(E)>0}}\, \mathrm{b}\,(\mathrm{T_i})\mathrm{ch}(\mathrm{A}|\mathrm{E})$

Because GPP places no direct constraints on an agent's conditional credences, it's compatible with GPP that an agent have conditional credence $\mathrm{b}(\mathrm{A}|\mathrm{T_i}) = 0$ when $\mathrm{ch_i}(\mathrm{A}) \neq 0$. However GPP has other unpalatable consequences: Pettigrew (2013a) shows that GPP is inconsistent with Bayesian conditionalization, and that GPP requires us to assign credence zero to chance functions which are themselves modest but which assign a nonzero probability to immodest chance functions.

Because NP is the least internally problematic principle, and because it fits best with other expert principles, I will focus on it for the remainder of the paper. Now for some notes and refinements: for Lewis, the chance theory *T* consists of a set of history-to-chance conditionals[54]. We can think of this function as taking as its input an ordered pair of propositions: a description of the history of a world and a proposition which is not implied by that history. The output of the function is a real number between zero and one.

[54]This way of understanding Lewis' chance function follows from a substantive claim about which propositions are admissible, namely, all and only those about the past.

Before we move on, it's worth noting that it's not clear why the Lewisian chance function needs to be a probability function. Lewis' defense was simple: the Principal Principle is a constraint on rationality. It is irrational for agents to have nonprobabilistic credences. This would lead to a contradiction if the chances were not probabilities. Hence the chance function must be a probability function. But this justification is indirect, and simply stipulates that the PP (a) is a requirement of rationality, and (b) can be added to the other requirements of rationality without inducing inconsistency. As things stand on the Lewisian proposal, neither of these stipulations admits of justification[55].

Fit and Accuracy

We now have in sight the two sides to the chasm we wish to bridge. On the one side we have an explication of chance's status amongst the laws (§[subBSA]). On the other we have formal principal linking it to our credences (§[chance credence link]). The question we now wish to answer is: why does the chance theory which best balances fit with simplicity deserve our deference in the sense given by NP?

For the same reason, I hold, that we defer to those that we consider to be exerts. It's natural to count someone as an expert just in case she is maximally accurate; we should defer to another agent in some domain if and only if we take that agent's beliefs to be better approximations of the truth than ours. Similarly we should defer to the objective chances if and only if they are closer to the truth than our partial beliefs. Consequently, we should

---

[55]There have been numerous attempts to justify PP and NP, but none have been successful. For a thorough overview, see Strevens (1997).

take *fit*, the virtue which measures how informative a chance theory is, to be a measure of accuracy: closeness to the truth.

Measuring Accuracy

There is a deep and growing literature on accuracy measures for credence functions (Joyce (1998), Leitgeb and Pettigrew (2010a)). The idea is simple: we take one credal function to be maximally accurate at a world. Call this the *vindicated function*. If our aim is truth, we will take the vindicated function at a world *w* to be the truth-function at *w*, v_w, where the truth function assigns 1 to truths and 0 to all falsehoods.

We then devise a measure of *distance* between credal functions. The accuracy of a credal function then is its distance from the vindicated function. The distance measure is a function $D(b_i, b_j)$ which measures the distance between two credal functions. If both functions are the credences of agents, the distance measure will tell us by how much they disagree. The standard distance measure takes *distance* to be the sum of the squared differences between the credences of the agents in each proposition:

DISTANCE: $D(b_i, b_j) = \Sigma_{A \in F}(b_i(A) - b_j(A))^2$,

where we assume for simplicity that the two agents have credences in the same (finite) set of propositions, *F*.

Now that we have a maximally accurate function v_w and a distance measure, we have all we need to compare the accuracy of partial beliefs at a world *w*. For we can define the *in*accuracy of a credal function *b*(*|*) as distance from maximal accuracy:

INNACCURACY: $I(b, w) = \Sigma_{A \in F}(b(A) - v_w(A))^2$

Now that we have in hand a notion which will allow us to compare the accuracy of credal functions, we can state the thesis that the fit of a chance theory varies with respect to its accuracy; chance theory *T* has a higher *fit* score than *T** at a world *w* if and only if ch_T is less inaccurate than ch_{T^*}.

FIT AS ACCURACY: $fit(T, w) > fit(T^*, w) \equiv I(ch_T, w) < I(ch_{T^*}, w)$

I will not argue directly for FIT AS ACCURACY. Rather, I will show how *fit as accuracy* provides a better explanation of the normative force of the NP, more directly explains the fact that the chances are probabilities, and reject objections to the claim. FIT AS ACCURACY does better than traditional measures of fit in some respects and worse in none[56].

Fit, Accuracy, and the BSA

Minimizing inaccuracy is a good first step for out account of the chance function. But it should be clear that the chance function cannot merely be that ur-credence function which minimizes inaccuracy. For we know what this function is: it's the truth-function. And we know that the truth function is not the chance function.

---

[56]Here may be the best time to address a worry the reader may have: how does fit as accuracy deal with worlds, like ours, which have infinitely many chancy events? At least as well as traditional measures of fit. Like traditional measures of fit, we can trump up a finite set of 'test propositions', and measure the accuracy of the chance function in terms of these, as suggested in Elga (2004). If we have a countable, ordered set of events–or countable, ordered set of test propositions (as discussed below), we can also take the inaccuracy of the system to be the limit $\lim_{n\to\infty} \Sigma_{A_1, A_n}(b(A_i) - v_w(A_i))^2$.

The BSA, of course, has a response to this: fit trades off against simplicity. But Lewis is unclear about how one chance function can be simpler than another, as I discussed in §[BSA and PP].

Although Lewis' views on this subject aren't clear, two options are apparent in Lewis (1984). On the first, the simplicity of a chance function depends on the values of the chances, measured by the length of their decimal expansion. A chance function which assigns an event 0.5 rather than 0.4999 is more simple, on this view. The second takes chance functions to be sets of history-to-chance conditionals and measures the simplicity of the set of history-to-chance conditionals. Neither of these is satisfactory. The former is unsatisfying because it is hard to see why the difference between decimal expansions matter, the second is unsatisfying because time is dense. Any set of history-to-chance counterfactuals will thus be infinite–and so it's difficult to see how we will compare their simplicity.

Here I'll make two new proposals.

Simplicity: Conservatively Modified

The first involves replacing the history-to-chance conditionals with functions from fundamental quantities to chances. Developing this proposal requires a little groundwork: first, we should think of the Humean mosaic as composed of fundamental quantities; these quantities are determinables, and their values determinates. For example, *mass* is a fundamental quantitative determinable and *1g* is a determinate. Thinking of the mosaic this way has two advantages: first, it's closer to the way scientists think the world is[57]. We can allow the chances to be a function with a variable for each fundamental quantity, which takes as values determinates of this quantity. The chances, then, will be given by a function from quantity variables to probability-variable pairs. For example, quantum mechanics gives probabilities for some quantities (e.g., position) as a function of the values of other quantities (the amplitude of the wavefunction).

We can evaluate the simplicity of this function the same way we do that of deterministic laws: by looking at its syntactic complexity when it is stated using variables corresponding to the fundamental quantities. But we might want to add other requirements as well: for example, we may prefer chance functions which are continuous and so are such that small changes in values of our variables correspond to small changes in the chances. We may also prefer to rank the simplicity of such theories not (or not only) by their syntactic simplicity, but instead include *symmetry* considerations. A chance theory which respects spacial and temporal symmetries is much more practical than one which does not. If the chance of an outcome depends not just on local features of a chance setup but instead on the setup's location in spacetime, embedded agents (folks like us) would be unable to divine the correct chance setup by repeating qualitatively identical experiments[58].

---

[57]'Closer' because gauge freedom at the fundamental level suggest that the fundamental qualities of the world are more akin to graded relations than quantities. This debate is ongoing, and not one that I wish to engage here.

[58]This requirement is noted by Arntzenius and Hall (2003:179) who write that ``Your recipe for how total history determines chances should be sensitive to basic symmetries of time and space–so that if, for example, two processes going on in different regions of spacetime are exactly alike, your recipe assigns to their outcomes the same single-case chances." The difference between the view advocated by Arntzenius and Hall and the view

Simplicity: Radically Modified

My second proposal follows Hoefer (2007) and Ismael (2013) in taking conditional probabilities to primitive, where '$ch(A|E) = x$' means that the chance that an event of type E is of type A is x. We can then take the simplicity of the chance function to depend only on how many different situation-types it distinguishes between.

We can understand this typing extensionally by taking each type to be a set of events, so that x is of type E just in case $x \in E$. Each chance function will partition the events of the world into types, the set of which, T, is such that $E \in T$; we can then compare the simplicity of chance theories by comparing their typing schemes. If all of the types $E \in T$ are subsets of the types $E^* \in T^*$, then T^* is a simpler typing scheme than T, and a chance theory C based on T^* is simpler than a chance theory C* based on T.

Both proposals allow us to take the simplicity consideration to be tied to our evidence for the chances. For the simpler a chance theory is, on either measure, the more opportunities we have to observe it. The connection between frequency and chance goes in two directions: we gain information about the frequencies from the chances, and we gain information about the chances from the frequencies. In order for the chances to be epistemically accessible to us, we need to be able to infer them from observation. And in order to observe them, we need a broad class of events whose outcomes are assigned the same chance. So simpler chance theories are more *epistemically accessible*.

---

here advocated is that they take respecting spaciotemporal symmetries to be a requirement on chance functions rather than merely a good making feature.

Simplicity and Evidence

Our chance function should deliver to us the most accurate beliefs available to us given our evidence. This requires the chance function to be invariant over a broad enough set of events for us to observe the frequency of outcomes for those event types. It also requires the chance function to yield the same chances for any two situations which our evidence does not permit us to distinguish between. If, prior to performing observing an outcome of a chance setup, E and E* cannot be distinguished, then $\mathrm{ch}(\mathrm{A}|\mathrm{E}) = \mathrm{ch}(\mathrm{A}|\mathrm{E}^*)$. This gives us an absolute lower bound for the simplicity of our chance theories: they must respect qualitative indistinguishability:

QUALITATIVE INDISTINGUISHABILITY: For all *A*, *E*, and *E**, If *E* and *E** describe qualitatively indistinguishable situations, then $\mathrm{ch}(\mathrm{A}|\mathrm{E}) = \mathrm{ch}(\mathrm{A}|\mathrm{E}^*)$.

The point of QUALITATIVE INDISTINGUISHABILITY is simple: our chances cannot distinguish between situation which we cannot distinguish between. The chances need to be informative for agents like us, who have only limited, local access to information about the world. Since we want the chances to be the most accurate credal function *we have access to*, the chances can only distinguish between those situation we can distinguish between.

QUALITATIVE INDISTINGUISHABILITY may at first seem circular: our notion of what counts as distinguishable for what is tied to our notion of laws and objective chances. How, then, can we determine whether two setups are indistinguishable before we know what the laws and objective chances are? The worry, then, is that this requirement is toothless. What it is, we may think, for two setups *E* and *E** to be *distinguishable* is for $\mathrm{ch}(\mathrm{A}|\mathrm{E}) \neq \mathrm{ch}(\mathrm{A}|\mathrm{E}^*)$. This worrisome thought is mistaken.

First, it is false that we have no notion of what counts as indistinguishable from what prior to our account of laws and chances. For we don't yet have the final theory of laws and chances, but we do have a lot of true beliefs as to which situations are qualitatively indistinguishable. Two double-slit experimental setups are qualitatively indistinguishable, provided they're made of the same materials and are the same size, even if they are in different laboratories. Two shuffled decks of cards are qualitatively indistinguishable before the first card is drawn.

Second, QUALITATIVE INDISTINGUISHABILITY puts a limit on the grain a chance-theory can have. Because if two situations are qualitatively identical, they are clearly qualitatively indistinguishable. So our chance theory cannot assign different single-case chances to intrinsically identical events without violating QUALITATIVE INDISTINGUISHABILITY. But QUALITATIVE INDISTINGUISHABILITY is broader: two situations can be indistinguishable without being identical.

Third, even as an *internal* requirement on packages of law and chance, this requirement has teeth. For it requires situations to differ in *more than the chance of their outcomes* according to the theory. For a two setups to be assigned different outcome chances, there must be some *other* difference between them, either in terms of their internal distribution of fundamental properties or their causal history.

To sum up: I hold that chance functions fit the world better when they are more accurate, as measured by the Brier score. But the chances are not maximally accurate because they must respect two evidentialist constraints: first, we must be able to gain evidence for them by observing frequencies. Second, they cannot make distinctions between events which exceed our ability to distinguish between those events. The chance function, then, is the most accurate credal function for which we can gain evidence through observation of frequencies and employ to constrain our credences about future events.

Extensions to the Infinite

Accuracy measures are typically taken over a finite set of propositions; this may[59] be reasonable if the propositions are meant to represent a finite agent's credal state. But the propositions discussed here represent the events at a world like ours–a world with infinitely, and probably a continuum, of events. Fortunately the measure of fit advanced here does no worse, and may do better, than the traditional measure in infinite worlds.

Lewis' measure of fit notoriously runs into problems in worlds with infinitely many events. Each such world is assigned the same probability by each chance system: 0. So each chance system gets the same minimal score for fit. Lewis thought that this could be solved by introducing infinitesimal probabilities (Lewis, 1994), but Adam Elga has shown that this is hopeless (Elga, 2004). Elga offers a suggestion: the Lewisian should take the fit of a world to be measured NOT by the probability given assigned to the world–which here we can think of as the conjunction of all true propositions–but instead by the probability assigned to *special* a set of *test propositions*. Elga further suggests that the Humean appeal to those test propositions discussed in (Gaifman and Snir, 1982), who discuss usefulness of taking probabilities to be a function of structured rather than unstructured propositions. This tactic is easily co-opted by the advocate of *fit as accuracy*. Here, though, I'd like to suggest some alternatives:

[59]Or may not.

Gaifman and Snir's formal system involves a predicate calculus, the names of which are refer to events which are countably infinite but ordered. The predicates can be interpreted as event types. The idea is that each world consists in an infinite progression of discrete events which fall into a finite class of types. A simple interpretation describes a world of coinflips. The names (*'1', '2', ..., n, ...*) refer to flips of the coin, and the predicates *H, T* refer to the outcomes heads and tails. They then show how to define conjunctive propositions which, at worlds which are 'normal' according to a probability distribution, are given probability 1. These are Elga's test propositions.

If the proponent of FIT AS ACCURACY follows Gaifman and Snir, she will similarly count chance systems as maximally accurate if they give Elga's test propositions probability 1 and maximally innaccurate if they assign them zero. But if she can appeal to the same formal language involving ordered events, she can define the fit of a system at an infinite world directly as the limit of a normalized Brier score:

$$\mathrm{I}(\mathrm{ch}, \mathrm{w}) = \lim_{\mathrm{n}\to\infty} \frac{1}{\mathrm{n}} \Sigma_{\mathrm{A}_1,\mathrm{A}_\mathrm{n}}(\,\mathrm{ch}(\mathrm{A}_\mathrm{i}) - \mathrm{v}_\mathrm{w}(\mathrm{A}_\mathrm{i}))^2$$

But she need not be wedded to this proposal. Instead, she is welcome to take the fit of a system to be given by its accuracy in terms of a finite set of test propositions whose members are determined by our epistemic goals. Our motivation for finding laws and objective probabilities is to extend our knowledge from the unobserved to the observed; we might take the test propositions to be those which are easiest to observe or those which we are most interested in predicting. I am not convinced that we can determine which these are without examining particular physical theories; in fact, which events are described by our test propositions may be irreducibly theory dependent (though constrained by EVIDENTIAL EQUIVALENCE).

This freedom in choosing test propositions gives us another avenue to move the best system closer to our goals as embedded agents. The test propositions described by Gaifman and Snir concern the totality of events; while we might gain information about these propositions by watching some local things happen, these propositions don't describe events that we can observe directly. We may instead prefer a set of test propositions which concern local goings on, and thereby increase the epistemic accessibility of the laws. We learn the laws by observing local events; we apply them by predicting local events. Why then should we take the lawmaking facts to be global events?

For example, on the Bohmian account of quantum mechanics, probabilities are given by an initial probability distribution working in tandem with a global wavefunction. However, when systems are sufficiently isolated, they can be described using a *local* wavefunction, which we will call an "effective" wavefunction (Goldstein et al). When such isolated systems interact with their environments, they undergo an apparent collapse; the probabilities of different post-collapse states are given by the Born rule. Apparent collapses provide an ideal set of test propositions: if there are finitely many particles, there may be finitely or countably many of them, the theory provides clear probabilities for each post-collapse outcome, and apparent collapses can be easily observed by embedded agents, who can then evaluate the theory by its accuracy in predicting post-collapse states. According to Ghirardi-Rimini-Weber quantum mechanics, these collapses are not merely apparent; in both cases the theory provides us with a finite or countably infinite set of events to which we match its frequencies[60]. Similarly, the Cosmological Measure Problem is the problem of taking an unaccountably infinite panoply of universes and finding a finite or countable set of events to which we can match our probabilities.

Chance and Epistemic Utility

This definition of chance fits neatly into an argument for the New Principle given by Pettigrew (2012). In §[epistemic utility and NP] I will present Pettigrew's (2012) proof. In §[vindicating chance] I will argue that the account of chance I have provided can be used to underwrite a key premise in Pettigrew's argument. In §[objections considered], I'll consider some objections to this argument, including those which led Pettigrew himself to abandon it in Pettigrew (2013).

An Epistemic Utility Argument for the Principal Principle

Pettigrew's (2012) proof is based on the notion that chance is to credence as truth is to full belief. Following Hajek (MS), Pettigrew holds that our credences should aim at the chances just as our beliefs should aim at the truth. To represent this goal, Pettigrew introduces the *chance-based Brier score* as a measure of epistemic utility:

CBBS: $\mathrm{I}(\mathrm{b},\mathrm{w}) = \Sigma_{\mathrm{A}\in\mathrm{F}}(\,\mathrm{b}(\mathrm{A}) - \mathrm{ch}_{\mathrm{w}}(\mathrm{A}|\mathrm{E}))^2$,

where *E* is the total evidence admissible to an agent at a time. We will say that CBBS measures epistemic utility in the presence of *E*. Taking *cbbs* as a measure of epistemic utility follows from the claims that (1) the ideal credence function at *w* is $\mathrm{ch}_{\mathrm{w}}(* \mid *)$ and (2)

[60]Provided that there are finitely many particles; if there are infinitely many particles applying any of these proposals (or any of these theories) is extremely difficult.

the *dis*utility of a credal function is proportional to its distance from the ideal function, as measured by the sum of squared differences. Note that the CBBS is just INACCURACY from §[Fit and Accuracy] with the 'vindicated' function taken to be the chance function rather than the truth function.

The remaining premises in Pettigrew's argument are imported from Joyce (1998). First, we assume (3) *dominance*: if credence function *b* has a higher epistemic utility than *b** at all worlds, and there is no other credence function that has a higher epistemic utility than *b* at all worlds, then it is irrational for an agent to employ *b**.

Pettigrew shows that (1), (2), and (3) together imply that agents are irrational to adopt credal functions which do not obey the axioms of probability and the Principal Principle– and that by slightly tweaking CBBS we can show that agents must obey the New Principle. Pettigrew's proof relies on taking the chance function to be a probability function.

Pettigrew provides little support for claim (1), that the ideal credence function at *w* in the presence of E is $\mathrm{ch_w}(* \mid \mathrm{E})$. His argument rests on the claim that chance is to credence as truth is to full belief, which Al Hajek has defended in unpublished work. But it is difficult to see how this claim could be defended without a metaphysical account of chance. We take truth the be the aim of belief not as a basic posit but because it comports well with our theories of truth. On the most naïve correspondence theory, the true propositions represent actual the world. Our beliefs are those of our mental representations which, ideally, match the world. This world-dependence is what makes truth at appropriate aim for belief at @; in order to accept the claim that our credences should aim at the chances, we need a similar account of the world-dependence of the chances.

Pettigrew also relies on but fails to support the claim that the chances must be probabilities. To my knowledge, there are two defenses of the claim that the chances must be probabilities. The first relies on frequentist or hypothetical frequentist accounts of chance: frequencies, as ratios of outcomes, are probabilities, so the chances must be probabilities. But this justification relies on a false theory of chance, as discussed in §[subBSA]. The second justification goes *via* the Principal Principle. But in the context of Pettigrew's proof, this is circular. A new vindication is given below.

Vindicating Chance

If the chances obey our two constraints, we have an argument for the claim that the chance function is the ideal credence function and a reason to believe that the chance function is a probability function. I'll start by arguing that it is irrational to fail to respect the chance function arrived at by taking the lower bound of simplicity, that is, the least simple chance function compatible with EVIDENTIAL EQUIVALENCE.

Recall that the chances, according to FIT AS ACCURACY and QUALITATIVE INDISTINGUISHABILITY, are the credence function that maximizes accuracy while treating evidentially indistinguishable setups as equivalent. The chance function, then, is the most accurate credence function *which obeys the same evidence-based constraints that we do*. In this sense, then, chance is to credence not as truth is to belief, but as *knowledge* is to belief[61].

Now suppose that someone knowingly fails to match their credences to the chances. Then she either employing a credence function which fails to respect QUALITATIVE

[61]Or perhaps as *justified belief* is to belief.

INDISTINGUISHABILITY or she is employing a credence function which respects QUALITATIVE INDISTINGUISHABILITY but is less accurate at her world than another such function.

If she fails to respect QUALITATIVE INDISTINGUISHABILITY, then her credences in the outcomes of chance setups depend on more than her evidence; while she cannot distinguish between two situations, she has different attitudes towards their outcomes. Her beliefs overshoot her evidence; she admits that she has no way to tell one situation from the other, but nonetheless assigns them different confidences. Her difference in attitude cannot be based on any evidence; this is irrational[62].

And if she respects QUALITATIVE INDISTINGUISHABILITY but is using a less than maximally accurate credal function, she is also irrational–because she is knowingly employing a credence function farther from the truth than she could. Hence if she fails to respect the chances she is irrational.

This gives us reason to believe that our credences should aim, not directly at the truth, but instead at the objective probabilities. For just as our goal for full belief is not merely to have true beliefs, but to have knowledge, our aim for partial belief should not be merely to have accurate beliefs, but to have accurate and well-supported beliefs. Since our credences cannot be more accurate (while retaining evidential support) than the chance function, we should take the objective chances at our world to be the target of our credences.

---

[62]I am not advocating an indifference principal here. An indifference principle requires one to have *equal credence* when one has no evidence bearing on a situation. I frankly do not care what credence one has in different situations; it's compatible with all I have said that you be as confident as you like that your coin will come up heads (or have as wide a spread of credences as you like); just have the same confidence each time you flip it, if you can't tell the flips apart.

Finally, this account of accuracy puts the claim that chances are probability functions on more solid ground. Previous attempts to justify this claim have gone either by way of linking chances directly to relative frequencies, which are ratios of integers and guaranteed to obey the axioms of probabilities, or indirectly through the Principal Principle. The former method is misguided as chances are not relative frequencies, actual or hypothetical. The second method is more subtle, but still not ideal. According to this justification, the Principal Principle is a norm of rationality. It is irrational to have credences which are not probabilities. So the chances must be probabilities.

This second justification undercuts our ability to justify the Principal Principle by appealing to the chances. For take some purported justification of the Principal Principle in terms of our metaphysics of chance; will this appeal to the fact that the chances are probabilities (as Pettigrew's does)? If so, the 'justification' is circular. But suppose the justification does not take this as an assumption. How then can we be assured that the Principal Principle is consistent with other constraints of rationality for partial belief? And if we are not sure of this, how can we accept the justification?

By making accuracy a constitutive feature of chance, we are able to answer this question. The same constraints that require our credences to be probabilities require the chances to be probabilities: namely, any non-probabilistic chance function will be accuracy dominated by a probabilistic one (see Joyce (1998)). But since our chances must be the most accurate credal function at some grain of simplicity, they will not be accuracy dominated. So they will be probabilities.

Objections Considered

Pettigrew's Objection

Pettigrew (2013) rejects CBBS and (1) because they conflict with two intuitive epistemic principles:

EPISTEMIC UTILITY OF FULL BELIEF: an agent at a world *w* has maximal epistemic utility only if she has full beliefs in all propositions which are true at *w* and full disbelief in all propositions which are false at *w*.

EPISTEMIC UTILITY OF FULL BELIEF is meant to follow from truth as the aim of full belief.

PARTIAL LOCKEAN THESIS: if an agent has full belief in a proposition, then she must have credence of greater than 0.5 in that proposition.

Taking the chances to be the ideal credal function requires agents to often have less than .5 credence in truths, and so (if the PARTIAL LOCKEAN THESIS holds) requires them to disbelieve truths. But then an agent with maximal epistemic utility for her partial beliefs could not have maximal epistemic utility for her full beliefs, given EPISTEMIC UTILITY FOR FULL BELIEF.

I am not convinced that PARTIAL LOCKEAN THESIS holds. For imagine we adopt a simple functionalist picture of belief, where to believe that *if P then Q* is to be such that if one desires that *Q* one acts to bring about *P*. Then for a utility maximizer to believe *if P then Q* she merely need to have a higher credence in *if P then Q* than its alternatives. But if *if P then Q* has more than one alternative, she may well have a credence of less than .5 in it while still being more confident in it than its alternatives.

But we can shelve these doubts about the PARTIAL LOCKEAN THESIS and focus on EPISTEMIC UTILITY FOR FULL BELIEF. For EPISTEMIC UTILITY FOR FULL BELIEF encourages agents to disregard their evidence; while an agent will be correct in one sense if she has full belief in all truths and full disbelief in all falsehoods, there is another sense in which she will clearly be misguided: her beliefs are not properly founded on evidence. CBBS requires us to reject EPISTEMIC UTILITY FOR FULL BELIEF, not because we do not value the truth, but because we do not think that merely believing truths is valuable. Rather, *knowing* truths is epistemically valuable; similarly, having *well-founded* accurate credences is valuable.

Hall's Objection

Hall (2003) provides a positive argument against reductionist proofs of chance-credence principles. Suppose you live in a world with chancy laws governing collisions between particles. All collisions are either elastic or inelastic, and the best systemetization of these collisions gives a probability for the collision's type which depends only on the mass of the particles. Suppose further that there is at least one collision such that there are no other collisions between particles with that mass–a unique collision:

``...you are about to observe a unique collision. The sort of categorical information that the reductionist says you possess, in virtue of knowing what the chances are, is therefore *simply not relevant*. All that is relevant is that you are about to observe a collision of a certain type, and the frequency of elastic outcomes among collisions of that type is either 1 or 0. If anything, it seems that the principle of indifference might apply–in which case C(A)=0.5" (Hall, 2003:109)

Hall's argument begs the question. Every collision is in some sense unique: each collision happens at a different spaciotemporal location. And each collision is similar, purely in virtue of being a collision. *What counts as evidence depends on which laws best unify the phenomena at a world.* The claim that the ``sort of categorical information you posses in virtue of knowing what the chances are, is therefore *simply not relevant*" is question begging–which bits of categorical information are relevant to our credences depends on the laws. The chances are justified by (a) their accuracy, and (b) what is suitable to count as evidence for some outcome: namely, the most similar events at the world.

Conclusion

Lewis proposed a humean account of chance which sought to directly link the chances to the frequencies at a world. I've provided a similar account which ties the probabilities not directly to frequencies, but instead to the *accuracy* of a chance function. I've also provided a refined account of the *simplicity* of a chance function, which more directly links simplicity to our evidence for a chance theory. I've argued that this account of chance underwrites a proof of the PP.

Conclusion

This dissertation is based on the old insight that the laws of nature are at their roots a description of the world rather than a ruler of it, together with the new insight that they are a description that *agents like us* can uncover. This insight allows humeans to overcome a number of interesting puzzles that besets their view: why are scientists routinely interested in dynamic rather than static information? Why do we have a number of independently operating scientific disciplines all describing the same world? Why should we pay attention to the objective chances?

But many puzzles remain. In this chapter, I'll offer some suggestions about where this view can go. First, the view of objective chance defended here can be refined and extended. Second, humeanism faces challenges in accounting for some of the bizarre consequences of quantum mechanics. Finally, humeanism has notorious difficulties in dealing with certain counterfactuals. Here I'll briefly discuss the way I see forward on these problems.

Multiverse theories in cosmology and the many worlds account of quantum mechanics both make an apparently irreducible appeal to *de se* probabilities: their predictions come via a probability that *you* will see an event rather than a chance that the event will occur. I believe that the humean of objective chance I defend here may help shed light on how irreducibly *de se* probabilities can feature in physical theories: the account of chance offered here takes chances to be primarily informational tools for agents. Nothing prevents the relevant information from being *de se* rather than *de dicto*. Both typicality assumptions and *de se* probability feature ineliminably in fine-tuning arguments. These arguments rely crucially on the assumption that our universe is atypical; when these arguments take place with the paradigm of cosmic inflation in the background, *de se* probabilities are used to ground these typicality assumptions. But in physics and philosophy of physics our understanding of what's required for a typicality assumption to be justified and our picture of *de se* probabilities are both extremely limited. In order for us to properly evaluate these arguments we need a deeper understanding of both.

First, this dissertation worked started from the assumption that humeanism *works*–that is, that armed only with a categorical mosaic of properties, we can give satisfying accounts of nomic necessity, counterfactuals, causation, and explanation. But a number of authors have challenged the humean project's very foundation. I'll briefly discuss the two of these challenges that I find most troubling, and make some noise in the direction of solving them.

The first such challenge argues that we have evidence from physics that the world, at its most fundamental level, does not consist of a mosaic of properties, localized at spacetime points, and unconnected by any necessary relations. The first bit of physical evidence for this arises from quantum mechanics. When two particles are entangled at the quantum level, their combined state is not a function of their individual states taken together; the state of the two-particle system is something over and above the states of the two particles individually. The second bit comes from quantum field theory; there, the basic qualitative items seem not to be properties but fiber bundles, and their structure is very unlike that of the mosaic of unconnected properties envisioned by traditional humeans.

There are two extant humean responses to this worry. The first, due to Loewer (1996), takes the fundamental space not to be the 4-dimensional world of spacetime, but the much higher-dimensional space in which the quantum wavefunction operates. This view, a variety of wave-function realism, notes that *if* the wavefunction is real, then we are stuck with its high-dimensional space anyway; to paraphrase Jaden Smith, how can the wave be real if its space isn't real?

The second response argues that the wavefunction is *not* real. this view, advocated by Miller (2014), Perry and Bhogal (MS), and Chen (MS) holds that, at least on some interpretations[63] of quantum mechanics, the fundamental structure of reality consists of particles and locations. The wavefunction, on this view, is just a handy way of coding up information about the locations of the particles. Particle entanglement is a feature of the wavefunction, not of the particles themselves, and so the fundamental ontology of quantum mechanics is compatible with humeanism.

These responses together negate this worry about humeanism. We must either be realists about the wavefunction, or not. If we must be realists, then we should follow Loewer in taking the statespace over which it is defied to be the space over which the humean mosaic is spread. If we must be antirealists, we should follow Miller, Chen, and Perry and Bhogal in seeing the wavefunction as a bearer of information and nothing more.

If we take the latter route, we are presented with a more austere humeanism even then Lewis'. For on this view, all there is to the world is particles and their positions. All of their properties, coded into the wavefunction, are less than real, and instead only serve as markers to help us predict their future positions. This latter account of the wavefunction fits nicely with the epistomology-first sort of humeanism I favor. If the laws are primarily aimed at bringing us information about the world's fundamental level, it should not

[63]The Bohmian one.

surprise us if they introduce nonfundamental conceptual machinery to do so. However, the choice of concepts is not unlimited: for if we can use any predicate we want, we will find ourselves again with Lewis' dreaded predicate F (from §[BSA]). I believe that the account of the role of laws offered here, and the new virtues it introduces, can allow us to differentiate the acceptable forays into the nonfundamental from those unfit to feature in our most basic science.

Finally, humeanism cannot easily explain the fact that the laws could apply to very simple worlds. For example, quantum mechanics can be used to describe a world consisting only of a single particle in an infinite square well; Newtonian mechanics can describe a world consisting only of a single particle moving at one meter per second. But according to humeanism–even of the variety offered here–the laws at those worlds would not be quantum or classical mechanics; rather, the simplest package of information about those worlds would consist in a short summary of the behavior of the single particle. But this seems to get both the facts and counterfacts wrong: there could be just a single particle in worlds governed by these laws.

Here the humean is in a bind. But I see light in the view; for while the worlds described are *metaphysically* impossible, by the humean's lights, they are *epistemically* possible. And the humean holds that the laws are an epistemic tool, not a metaphysical one. So perhaps science is concerned not with metaphysical but rather with epistemic possibility. If so, these worlds are in fact possible in the sense relevant for science.

This response needs to be worked out in much more detail. But it provides hope that the humean can respond to even the most intractable charges leveled against her. And as humeanism offers us our best chance of explaining our interest modal facts, rather than asserting it, we should continue striving towards this explanation so long as hope lasts.

References

Albert, David (2000). *Time and Chance*. Cambridge, MA: Harvard University Press.

Armstrong, David (1983). *What Is a Law of Nature?* Cambridge: Cambridge University Press.

Armstrong, David (1997). *A World of States of Affairs*. Cambridge: Cambridge University Press.

Baumann, Peter (2005). ``Theory Choice and the Intransitivity of 'Is a Better Theory Than.' *Philosophy of Science* 72 (1):231-241.

Beatty, John (1995). ``The Evolutionary Contingency Thesis." in Wolters, Gereon, and James G. Lennox, eds, *Concepts, Theories, and Rationality in the Biological Sciences*. Pittsburgh: University of Pittsburgh Press.

Beebee, Helen (2000). ``The Non-Governing Conception of the Laws of Nature." *Philosophy and Phenomenological Research* 61: 571–594.

Beebee, Helen (2006). ``Does Anything Hold the Universe Together?" *Synthese* 149 (3):509–533.

Beebee, Helen (2011). ``Necessary Connections and the Problem of Induction." *Nous* 45 (3): 504–527.

Bird, Alexander (2007). *Nature's Metaphysics: Laws and Properties*. Oxford: Oxford University Press.

Callender, Craig, and Jonathan Cohen (2009). "A Better Best System Account of Lawhood." *Philosophical Studies* 145 (1): 1–34.

Callender, Craig, and Jonathan Cohen (2010). "Special Science, Conspiracy, and the Better Best System Account of Lawhood." *Erkenntnis* 73: 427–447.

Cartwright, Nancy (1983). *How the Laws of Physics Lie.* Oxford: Oxford University Press.

Cartwright, Nancy (1999). *The Dappled World: A Study at the Boundaries of Science*. Cambridge: Cambridge University Press.

Cross, Troy (2012). "Goodbye, Humean Supervenience." *Oxford Studies in Metaphysics* 7: 129–153.

Demarest, Heather (forthcoming). "Powerful Properties, Powerless Laws." in Jonathan Jacobs (ed.) *Putting Powers to Work: Causal Powers in Contemporary Metaphysics*. Oxford: Oxford University Press.

Dretske, Fred (1977). "What Is a Law of Nature?" *Philosophy of Science* 44: 248–268.

Earman, John, and John Roberts (2005a). “Contact with the Nomic: A Challenge for Deniers of Humean Supervenience about Laws of Nature (Part I).” *Philosophy and Phenomenological Research* 71 (2): 1–22.

Earman, John, and John Roberts (2005b). “Contact with the Nomic: A Challenge for Deniers of Humean Supervenience about Laws of Nature (Part II).” *Philosophy and Phenomenological Research* 71 (2): 253-268.

Eddon, Maya, and Chris Meacham (forthcoming). ``No Work for a Theory of Universals.” In Jonathan Schaffer and Barry Loewer (eds.), *A Companion to David Lewis*.

Ellis, Brian D. (2001). *Scientific Essentialism*. Cambridge: Cambridge University Press.

Elga, Adam (2004). “Infinitesimal Chances and the Laws of Nature.” *Australasian Journal of Philosophy* 82 (1): 67–76.

Forster, Malcom and Elliot Sober (1994). ``How to Tell When Simpler, More Unified, or Less ad hoc Theories Will Provide More Accurate Predictions.” *British Journal for the Philosophy of Science* 45 (1): 1-35.

Gaifman, Haim, and Marc Snir (1982). ``Probabilities Over Rich Languages.” *Journal of Symbolic Logic* 47(3): 495-548.

Goodman, Nelson (1955). *Fact, Fiction, Forecast*. New York: Harvard University Press.

Hall, Ned (2004). “Two Mistakes about Credence and Chance.” *Australasian Journal of Philosophy* 82(1): 93-111.

Hall, Ned (MS). “Humean Reductionism about Laws of Nature.” <philpapers.org/rec/HALHRA> accessed September 13, 2013.

Hicks, Michael and Jonathan Schaffer (MS) “Derivative Properties in Fundamental Laws.”

Hoefer, Carl(2007). ``The Third Way on Objective Probability: A Skeptic’s Guide to Objective Chance.” *Mind* 116 (463): 549-596.

Ismael, Jenann (2009). ``Probability in Deterministic Physics.” *Journal of Philosophy* 86 (2): 89-109

Joyce, James (1998). ``A Non-Pragmatic Vindication of Probabilism" *Philosophy of Science* 65 (4):575-603.

Joyce, James (2007). “Epistemic Deference: The Case of Chance.” *Proceedings of the Aristotelian Society, Vol. CVII, Part 2.*

Kitcher, Philip (1989). “Explanatory Unification and the Causal Structure of the World”, in Kitcher, Philip, and Wesley Salmon (eds.), *Scientific Explanation*, 410–505. Minneapolis: University of Minnesota Press.

Kripke, Saul (1980). *Naming and Necessity*. New York: Harvard University Press.

Kukla, Andre (1995). ``Forster and Sober on the Curve Fitting Problem." British Journal for the Philosophy of Science 46 (2): 248-252.

Lange, Marc (2009). *Laws and Lawmakers: Science, Metaphysics, and the Laws of Nature.* Oxford: Oxford University Press.

Leitgeb, Hannes & Richard Pettigrew (2010). ``An Objective Justification of Bayesianism I: Measuring Inaccuracy." *Philosophy of Science* 77 (2):201-235.

Leitgeb, Hannes & Richard Pettigrew (2010). ``An Objective Justification of Bayesianism II: The Consequences of Minimizing Inaccuracy." *Philosophy of Science* 77 (2):236-272.

Lewis, David (1980). "A Subjectivist's Guide to Objective Chance." In Richard C. Jeffrey, ed., *Studies in Probability and Inductive Logic*. University of California Press.

Lewis, David (1983). "New Work for a Theory of Universals." In *Papers in Metaphysics and Epistemology.* Oxford: Oxford University Press.

Lewis, David (1994). "Humean Supervenience Debugged" *Mind* 103 (412): 473–490.

Loewer, Barry (1996). "Humean Supervenience." *Philosophical Topics* 24 (1): 101–127

Loewer, Barry (2007). "Laws and Natural Properties." *Philosophical Topics* 35 (1/2): 313–328.

Loewer, Barry (2008). "Why There Is Anything Except Physics." In Jakob Hohwy and Jesper Kallestrup, eds., *Being Reduced: New Essays on Reduction, Explanation, and Causation*. Oxford: Oxford University Press.

Loewer, Barry (2009). "Why Is There Anything Except Physics?" *Synthese* 170 (2): 217–233.

Loewer, Barry (2012). "Two Accounts of Laws and Time." *Philosophical Studies* 160 (1): 115–137.

Maudlin, Tim (2007). *The Metaphysics Within Physics*. Oxford: Oxford University Press.

McKenzie, Kerry (forthcoming). "No Categorical Terms: A Sketch for an Alternative Route to Humeanism about Fundamental Laws." In M. C. Galavotti, S. Hartmann, M. Weber, W. Gonzalez, D. Dieks, and T. Uebel (eds.) *New Directions in Philosophy of Science*. New York: Springer.

Meacham, Christopher J. G. (2010). ``Two Mistakes Regarding the Principal Principle." *British Journal for the Philosophy of Science* 61 (2):407-431.

Mill, John Stewart (1843). *A System of Logic, Ratiocinitive and Inductive.* Reprint. New York: Echo Library 2009.

Mitchell, Sandra (1997). "Pragmatic Laws." *Philosophy of Science* 64: S468–S479.

Mitchell, Sandra (2000). "Dimensions of Scientific Law." *Philosophy of Science* 67: 242–265.

Pettigrew, Richard (2012). ``Accuracy, Chance, and the Principal Principle." *Philosophical Review* 21:2 241-275.

Pettigrew, Richard (2013a). ``Accuracy and Evidence." *Dialectica* 67 (4):579-596.

Pettigrew, Richard (2013b). ``A New Epistemic Utility Argument for the Principal Principle." *Episteme* 10 (1):19-35.

Ramsey, Frank Plimpton (1927). "Laws and Causality." Reprinted in D. H. Mellor (ed.) *F. P. Ramsey: Philosophical Papers.* Cambridge: Cambridge University Press 1990.

Rosenkrantz, Roger (1977). *Inference, Method, and Decision: Towards a Bayesian Philosophy of Science*. New York: Reidel.

Shoemaker, Sidney (1980). ``Causality and Properties". In Peter van Inwagen (ed.), *Time and Cause*. D. Reidel. 109-35.

Schaffer, Jonathan (2003). ``Principled Chances" *British Journal for the Philosophy of Science* 54 (1):27-41.

Strevens, Michael (1999). ``Objective Probability as a Guide to the World." *Philosophical Studies* 95 (3):243-275.

Thau, Michael (1994). ``Undermining and Admissibility." *Mind* 103 (412): 491-504.

Tooley, Michael (1977) "The Nature of Laws." *Canadian Journal of Philosophy* 7(4): 667–98.

Weslake, Brad (2014). ``Statistical Mechanical Imperialism", in Alastair Wilson (Ed), *Chance and Temporal Asymmetry*, Oxford University Press, Oxford, pp. 241–257.

Woodward, James (2013). "Simplicity in the Best Systems Account of Laws of Nature." *British Journal for the Philosophy of Science* 0: 1–33.

Woodward, James (2013). "Laws, Causes, and Invariance." In Mumford, Stephen, and Matthew Tugby, (eds.) *Metaphysics and Science*, pp. 48–72. Oxford: Oxford University Press.