Dynamic Humeanism

Mike Townsen Hicks

August 12, 2016

Abstract

Humean accounts of laws of nature fail to distinguish between dynamic laws and static initial conditions. But this distinction plays a central role in scientific theorizing and explanation. I motivate the claim that this distinction should matter for the Humean, and show that current views lack the resources to explain it. I then develop a regularity theory which captures this distinction. My view takes empirical accessibility to be one of the primary features of laws, and I identify features laws must have to be empirically accessible. I then argue that laws with these features tend to be dynamic.

Contents

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 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 paper, 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 boundary conditions. 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 and thereby explain our preference for dynamic laws.

I will not address criticisms from proponents of more metaphysically robust accounts of laws. 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.

1 The Best System

The regularity theory holds that laws of nature are merely generalizations. The 'merely' 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 counterfactuals (Lange [2007]). These bits of metaphysical machinery underwrite the necessity of the laws. Although each of these metaphysical machines is distinct they are united in holding that facts

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

[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]).

about nomic necessity, or facts with modal implications 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.

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.

1.1 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 sparsest 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[4]. The BSA has changed slightly since Mill's 1843 explication: modern regularity theorists have

[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 in Ramsey [1929]).

[4] Lewis introduces the notion of fit to account for probabilistic lawsystems in his ([1980]); addressing the issues raised by fit (see (Lewis [1994]) and (Elga [2004]) for a few) would take more space than this paper allows, so I will not discuss it further.

backed off of Mill's claim 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 (Loewer [1996], [2007]; Beebee [2000], [2006]; Callender and Cohen [2009], [2010]; and 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[5]. 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.

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 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. Many philosophers are leery of this constraint: it seems to invoke a suspicious degree of unobservable structure. Even otherwise orthodox Humeans like Loewer ([2007]) and Callender and Cohen ([2009], [2010]) develop accounts meant to dispense with it. As a preview, one of the advantages of the view I offer here is that it de-emphasizes syntactic simplicity, and so is less constrained by metaphysical notions of naturalness.

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

- 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*.

[5]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 favouring any lawsystem.

- 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 must 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].

Given this conception of laws, the virtues of simplicity and strength are quite natural: organizations of knowledge are better when they are more organized—that is, simpler[6]—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 §3.

2 The Best Is Not Good Enough

The BSA holds that being a part of the simplest, strongest system couched in perfectly natural terms is necessary and sufficient for being a law of nature. This is false. In §2 I will present arguments against the BSA. The first (§2.1), the argument from boundary conditions, shows that the BSA lacks the resources to distinguish between laws and boundary conditions. The second, (§2.2) an argument from scientific practice, shows that the virtues articulated by the BSA are not those sought by scientists. Hence, the BSA's explanation of the role of laws cannot adequately explain the norms on theory choice scientists do respect. In §2.3 I provide a counterexample to the BSA illustrating these failures.

[6]An anonymous reviewer points out that simplicity—when understood as syntactic simplicity in the language of natural properties—does not necessarily track organization. I agree; this looks to me to be another area in which Lewis's version of the BSA does not live up to its own standards.

2.1 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. And examples from the history of science are also easy to come by: the fact that causally isolated regions of the early universe were in thermal equilibrium would have made a simple and extremely informative law, but cosmologists instead saw this as a serious and unexplained problem until it was accounted for dynamically by inflationary cosmological models that arose in the 1980s (a criticism along these lines can be found in Roberts [2009]).

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

> [N]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.

2.2 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 (Loewer [2007]), 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. Hall eschews direct discussion of strength and simplicity in favour 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.

Scientists are not willing to make tradeoffs 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. Moreover, Newton's theory is more accurate than Kepler's theory: it accounts for more of the perturbations of planetary motion and does so with less error. Instead of sacrificing strength for simplicity, scientists start by formulating the strongest and most accurate 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 (Roberts [2008] pp. 8–9; and Woodward [2013a])[7]).

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

[7] Thanks to an anonymous referee for discussion here

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. According to the BSA, simplicity considerations are intended to make the laws cognitively tractable.

But philosophical research focused on simplicity considerations in the sciences finds 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 once they are updated on evidence. Others (Forster and Sober [1994]) hold that preferring simpler theories guards against measurement error. A thorough review of philosophical accounts of simplicity can be found in (Huemer [2009])[8].

Rosenkrantz ([1977]) explicitly formulates an objective Bayesian model of scientific inference, similar to one appealed to by Henderson ([2013]). Both hold that before we obtain evidence we ought to have equal confidence in all theories we believe are possible. We can think of this as a sort of fair judgement axiom—if we start our inquiry with a higher credence in one theory than another, our inquiry is biased.

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. 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 which theory is true before we accumulate evidence. And unless our priors are absurdly biased towards the complicated theories, evidence gives a preferential boost to the simpler ones.

Conversely, 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

[8]In addition to those reasons for favouring simpler theories discussed below, Huemer discusses a metainductive argument for favouring simpler theories (suggested by Nolan ([1997] pp. 331–2)) and an argument which notes that complexity has a lower bound but no upper bound (found in Jeffries [1973] pp. 38–9); these reasons provide no more reason to sacrifice strength for simplicity than those I consider here. Thanks to an anonymous referee for discussion of this point.

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) (see Forster and Sober [1994], and Huemer [2009] for a thorough and accessible introduction to this strategy.).

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 favouring simpler theories is truth-conducive—either for our predictions or our belief in the relevant theory. When a more complex theory will tell us more, or tell us more accurately, simplicity doesn't seem to matter.

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 neither do we nor should we look for laws merely as organizational tools, as the BSA suggests.

2.3 An Illustrative Example

In §2.1 I argued that our world's best systematization contains statements which are not laws. In §2.2 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 virtue of simplicity and strength don't match our scientific and epistemic interests. This counterexample deals with a merely possible world. In this world, there is no single best system, but, I argue, there is a set of generalizations 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 modelled 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[9]:

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

The faster something moves, the more Taylor's force opposes that motion. But each subsequent coefficient $(a, ..., d, ...)$ is much smaller than the one before it. So at low speeds it can be approximated as follows: $f = -av$, at moderate

[9]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.

speeds approximated as $f = -av - bv^2$, and so on[10].

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 has no finite expression, any systematization which contains it will be maximally unsimple. So it will be no better with respect to the traditional measure of simplicity than a mere list of all facts[11]. Hence there will be no unique simplest, strongest systematization: the system containing f_T will be tied with other equally strong infinitary lists. 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, even when one of these lawbooks merely lists facts. But more importantly, the BSA is blind to the 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 of facts has none of these features. So the simplicity requirement of the BSA is not capturing those features of laws which make science possible.

3 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

[10]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, and so on. But this is not explicable on the traditional BSA.

[11]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 $a_1, a_2, ..., a_n, ...$, such that there is a force $f_T = -a_1v - a_2v^2 - ... - a_nv^n - ...$' This would allow us to state our laws finitely, but would 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.

new approach to the regularity theory. In §3.1 I argue that the BSA focuses on the output of scientific inquiry—predictions and explanations—and ignores the inputs—experimentation. In §3.2 I develop an account of laws which takes them to be the midpoint of inquiry, resting between induction and prediction. In §3.3 I make some suggestions about how this view accounts for scientific virtues, and in §3.4 I show how this view handles the problem cases outlined in §2.

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]), Callender and Cohen ([2009]), and Woodward ([2013a]) justly complain, the best systems account gives us no guidance in weighing the incommensurable virtues it identifies. The Epistemic Role Account put forward in §3.2, 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.

3.1 The Epistemic Criterion

The unifying theme of the difficulties presented in §2 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 marshalling 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] p. 187)

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[12], and it's easy to see why. Here's a clear takedown by Fred Dretske:

[12]It's not clear which philosophers have held the NEA. It's mentioned by Ramsey ([1927]), Goodman ([1955]), and Dretske ([1977]), though none of these authors accepts NEA or cites specific supporters. Whether Mill accepted the NEA or was merely describing inductive practice, I honestly haven't the foggiest: 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*.

> 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.[13] 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 explanations and perform manipulations. Scientists are looking for generalizations that can be known by observing a subset of their instances in controlled situations.[14]

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 of (non-exhaustive) inductive support.[15]

3.2 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 scientist 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 observa-

[13]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.

[14]This is also noted by Hoefer ([2007]), for whom simplicity is 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] p. 463).

[15]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]).

tion and experimentation[16]. Many of the most telling counterexamples to the BSA—such as generalizations stating the number of particles in the 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. 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 simplicity.

The input role of laws—that they must be suited to be inferred by observation and experimentation—does not appear in the BSA[17]. 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.

[16] A similar insight forms the basis of John T. Roberts' ([2008]) Measurability Account of Laws. Roberts takes the laws to be those principles which guarantee the reliability of measurements. My view differs from Roberts' in two principal ways: first, I hold that the application of laws in generating predictions and explanations is just as important to their counterfactual stability as the conditions required to discover them. Second, I take it to be incumbent on a view of laws to go some way towards explaining why we are interested in counterfactuals. While Roberts effectively explains the counterfactual stability of laws in terms of our interest in counterfactually stable modes of measurement, this is where his explanation stops; there is no indication of why we would require our measurements to be counterfactually rather than merely actually reliable.

[17] One might worry that, while this requirement does not appear explicitly in the BSA, it does appear implicitly. After all, how else will we discover the laws, if not through empirical inquiry? But though on the BSA the laws must be discovered if at all through observation and experiment, there is no guarantee that they be discoverable. I argue below that explicitly requiring that the laws be subject to repeated observation and experiment rules out many of the counterexamples to the BSA. The counterexamples to the BSA, I claim, cite conditions which cannot be directly discovered; they are either completely empirically inaccessible to agents like us, or can only be inferred indirectly, only after the laws have be discovered and with the help of the laws. Thanks to an anonymous reviewer for discussion of this point.

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, 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.

3.3 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.

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. So it might appear that any account of laws which appeals to a notion of isolation has fallen into circularity. Such is not the case: we can briefly sketch an account of what I take a quasi-isolated subsystem—hereafter 'QIS'—to be which does not presuppose any laws[18].

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[19].

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.[20]

When introducing a notion which will play a central role in a theory, it's

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

[19] 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.

[20] Spelling out precisely what 'almost true' should mean here would take us too far afield, especially given the wide literature on truthlikeness or verisimilitude. 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—1 kg of mass is closer to 2 kg than it is to 20 kg. 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.

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.

3.3.1 Strength and the ERA

There are two ways in which a lawbook can provide information about subsystems of the universe: it 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 'local strength', as it's closer to the Best System's measure of informativeness—although importantly different, as we will see.

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*.

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 exceedingly 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, because not every system is sufficiently isolated[21]: the Pluto–Sun system is not a QIS, as other objects in the Kuiper belt exert a strong gravitational influence on Pluto.[22]

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

[21] An anonymous referee helpfully pointed out that, because GR is a fully local theory, the Pluto–Sun system is at least apparently an exact QIS of GR. This is because, although perturbations in the Pluto–Sun system may ultimately be due to Neptune and Kuiper belt objects, these can be accounted for by adding local sourceless perturbations to a model of the system. This is true; but we might prefer a theory which adds in information about the possible sources of changes in the curvature of spacetime. Such a theory will not count the Pluto–Sun system as sufficiently isolated.

[22] Our interest in having broadly applicable laws is not limited to wanting laws which apply in many situations; we also want laws that apply in many situations which have some qualitative variety. This is because we want to apply the laws not only in making predictions, but also performing manipulations. So we want to know what will happen if we change some variables in our QIS. This requirement is omitted from the main text to cut down on the clauses in the final view, but will reappear in a subsequent footnote. Keep your eyes peeled!

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. But, as Earman ([1984]) points out, 'the practice of science speaks not in favor of strength *per se* but for strength in intended applications.'[23]

The information provided by a globally strong lawbook may only allow agents to make predictions if they know just where in the world they are; but agents are often operating with very limited knowledge about their position in space or time. So we should seek a lawbook that can be applied by agents with little or no global knowledge.

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 eight planets and many more asteroids and meteors within the gravitational influence of the sun. We can now provide a local measure of informativeness:

STRENGTH: A lawbook L is (locally) 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 local strength and breadth of a lawsystem are not syntactic features of the way a system of laws is stated. Moreover, they are not intrinsic features of the system, nor are they (like Lewis' notion of strength) constant across all worlds which the system describes. Rather, they are relations between a system and a world[24]. Consequently the primary features of a system which we weigh against one another when looking for laws do not depend at all on the way the laws are stated, and depend not just on the semantic content of those law statements; they also depend on how easily those statements can be discovered and applied in our particular world.

[23]Thanks to an anonymous reviewer for pointing me to this quote.

[24]Thanks to an anonymous referee for discussion of this point.

3.3.2 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 Coulomb's law (1785) was added to the classical lawbook. 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 large-scale structures of 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 independently of the others

Thus to discover and test our lawbook we need subsystems that are QISs of only 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 constants of l.

So the solar system is a QIS of a reduction of the laws of classical mechanics: it can be described without Coulomb's 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. Similarly, gravity is sufficiently weak that it can be ignored when doing terrestrial experiments on charged bodies; so we can discover Coulomb's law without looking at the Law of Universal Gravitation. Importantly, each of these laws has many QISs and applies to many subsystems, so each individually is broad. 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*.

3.3.3 Simplicity in the ERA

On the traditional BSA, the need for strength 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 §2.2 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, akaike, and the bayesian solution to the curve fitting problem—take favouring 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, syntactic simplicity combines with independently motivated constraints on inductive reasoning to encourage us to favour simpler theories.

Simpler theories, then, are favoured 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*.

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.

3.4 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, local strength, simplicity, and modularity. Consequently I tenuously characterize the ERA as follows:

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

This characterization is tenuous because I am not so confident that these exhaust the features that we require to discover laws. But this characterization is important: the failure of the BSA shows clearly that we must respect the needs of the embedded agents when developing our account of laws. But realism requires us to surgically remove the agent in our final characterization: our account should not be such that we can change the laws by changing our beliefs or epistemic goals, it should make sense of the fact that worlds without agents still have laws, and it should give us the capacity to explain our epistemological and practical interest in laws non-circularly. $\text{ERA}_{virtues}$ shows us that this can be done. The role-based account in defended in §3.2 can motivate a set of criteria for lawhood that makes no reference to agents, epistemic notions, or modally robust properties.

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

1. The BSA cannot differentiate laws from boundary conditions.

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.

Meanwhile, the ERA gives us a guide to which parts of our theories are laws and which parts are not. Is the mass of the electron fixed as a matter of law, or is it a mere fact that all electrons have a mass of $9 \times 10^{31} kg$? Is the Past Hypothesis—the claim that the universe started out in a low-entropy state (Albert [2003])—a law? On the ERA, this is not a question of the syntax of these claims as they are stated, nor is it merely a question of their overall informativeness. It's a question of their applicability to subsystems[25].

To see how this works, let's work through an example. Roberts ([2009] p. 17) argues that the BSA cannot correctly differentiate laws from initial conditions in, among other places, Newton's theory of planetary motion. Given Newton's law of universal gravitation, one can vastly increase the informativeness of the lawbook by adding to it the masses of the planets and the radii of their aphelions (and, perhaps, the angle between their positions at some time). If the BSA is correct, then, one would expect these facts to feature as laws in Newton's system. Call Newton's actual laws (the three laws of motion together with the law of universal gravitation) NEWTON and this expanded lawbook NEWTON$^+$. As Roberts argues, Newton clearly took NEWTON rather than NEWTON$^+$ to be the laws: he saw the balance between the masses and distances of the planets as a surprising thing that cried out for explanation, not a nomically necessary fact. So the BSA gives the wrong result here: it tells us that the masses and distances between the planets should be a matter of law, when intuitively they are a mere contingent facts—one which could, at least in principle, itself be explained.

I claim that the ERA gets the right result where the BSA does not. To see why, note that by including the specific masses of each planet, NEWTON$^+$ applies, as a whole, to far fewer subsystems than NEWTON. While NEWTON gives us strictly more information about the entire solar system, it contains

[25] I conjecture that on the account here offered, both are laws, but do not have space to make arguments for that claim. I think the view I defend here will be especially useful to friends of the Past Hypothesis, who need to say why that claim is a law, but other bits of information about the initial state of our world are not. On my view, the Past Hypothesis is a law because (nearly) every subsystem of the universe starts out in a low-entropy state; other facts about the initial state aren't similarly broad.

laws that do not apply to, for example, the Earth–Moon system or the Sun–Jupiter system or the system composed of Neptune and her moons. (Supposing that the additional laws of NEWTON$^+$ are existential statements, these laws are false of the Earth–Moon system: there is not in that subsystem an object with mass 6.39×10^23kg, the mass of Mars). Even worse, in our universe there are many planetary systems with many different sets of masses and interplanetary distances[26]. So NEWTON$^+$, while more globally informative, is only barely more locally informative and considerably less broad than NEWTON.

Recall that the Law of Universal Gravitation is $F_g = \frac{GMm}{r^2}$. One might want an explanation of the fact that G has its value as a matter of law, but M, m, and r do not. There is after all no syntactic difference between 'G' and 'M' here. The ERA's explanation of this is simple: while the masses and relative distances of objects in different subsystems of NEWTON may vary, the relationship between their masses and the forces they exert on one another does not. So by including the value of G in the its lawbook, NEWTON loses no breadth.

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[27][28].

2. 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 favour 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 favour a simpler lawbook over a more complex, but equally strong, lawbook). How modular must the

[26]This fact is not strictly relevant to an imagined assessment of Newtonian gravity, as conceived in the *Principia*, as Newton and his contemporaries were unaware of these other planetary systems. Fortunately, there are within this solar system sufficiently many QISs of NEWTON that are not also QISs of NEWTON$^+$ that we are able to motivate the correct result without looking beyond our home system.

[27]Thanks to an anonymous reviewer for discussion of this point.

[28]Remember footnote 22? We need to invoke it here do explain why some boundary conditions, like the large scale uniform distribution of matter in the universe, are not laws. It is because the QIS that these laws apply to don't vary sufficiently, and so their breadth is limited; they don't give us the kind of information we can use in explanations or manipulations (which require us to identify difference makers over and above facts). Thanks to John Roberts for raising this example.

laws be? Modular enough for us to discover the fundamental constants, and to bootstrap our way into discovering the whole book.

3. 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 than a list of facts, the ERA correctly rules that its statements are laws, and the list of facts is not. Such a lawbook provides the sort of information we want: it is more useful for modelling the sort of systems we use to gain evidence and in which we apply the laws, and this advantage is captured in the ERA by the fact that it is locally stronger and broader than a lawbook which merely lists facts.

4. The BSA rested on a suspicious notion of metaphysical naturalness.

Recall that the BSA invoked natural properties to explain why we don't invent a language, and a system, where every fact is a law. But the ERA takes the primary trade off to be between breadth and local strength. Neither of these is syntactic, so neither needs to specify a special class of predicates to do its work. Of course, modularity and simplicity on the ERA make direct references to the variables of a theory, and so sadly are syntactic and language dependent. But these virtues kick in only after breadth and local strength have winnowed the space of acceptable theories, and these virtues are more user directed, so plausibly are more sensitive to the predicates we find useful than those which are metaphysically bedrock.

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

4 Conclusion

The Best Systems Account of laws was designed to distinguish laws from merely accidental generalizations. I argued in §1 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 §2.

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.

More work needs to be done, of course: the laws are tied to explanation, causation, and judgements of epistemic safety. Explaining why the laws are uniquely suited to play these roles is a large project. But characterizing the laws correctly, and accurately identifying their role, is the first step in this project.

Acknowledgements

Thanks to Thomas Blanchard, Harjit Bhogal, Marco Dees, Heather Demarest, Alison Fernandes, Ned Hall, Barry Loewer, Zee Perry, John Roberts, Jonathan Schaffer and two anonymous referees for discussion and comments on drafts of this paper. Thanks also to the John Templeton Foundation for financial support.

References

Armstrong, D. [1983]: *What Is a Law of Nature?* Cambridge: Cambridge University Press.

Armstrong, D. [1997]: *A World of States of Affairs*, Cambridge: Cambridge University Press.

Albert, D. [2003]: *Time and Chance*, Cambridge: Harvard University Press.

Baumann, P. [2005]: 'Theory Choice and the Intransitivity of "Is a Better Theory Than"', *Philosophy of Science*, **72** (1), pp. 231–41.

Beebee, H. [2000]: 'The Non-Governing Conception of the Laws of Nature', *Philosophy and Phenomenological Research*, **61**, pp. 571–94.

Beebee, H. [2006]: 'Does Anything Hold the Universe Together?' *Synthese*, **149** (3), pp. 509–33.

Beebee, H. [2011]: 'Necessary Connections and the Problem of Induction', *Nous*, **45** (3), pp. 504–27.

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

Callender, C. and Cohen, J. [2009]: 'A Better Best System Account of Lawhood', *Philosophical Studies*, **145** (1), pp. 1–34.

Callender, C. and Cohen, J. [2010]: 'Special Science, Conspiracy, and the Better Best System Account of Lawhood', *Erkenntnis*, **73**, pp. 427–47.

Cartwright, N. [1983]: *How the Laws of Physics Lie*, Oxford: Oxford University Press.

Cartwright, N. [1999]: *The Dappled World: A Study at the Boundaries of Science*, Cambridge: Cambridge University Press.

Cross, T. [2012]: 'Goodbye, Humean Supervenience', *Oxford Studies in Metaphysics*, **7**, pp. 129–53.

Demarest, H. [forthcoming]: 'Powerful Properties, Powerless Laws', in J. Jacobs (*ed*), *Putting Powers to Work: Causal Powers in Contemporary Metaphysics*, Oxford: Oxford University Press.

Dretske, F. [1977]: 'What Is a Law of Nature?' *Philosophy of Science*, **44**, pp. 248–68.

Earman, J. [1984]: 'Laws of Nature: The Empiricist Challenge', in Bogdan, R. J. (*ed*), *D. M. Armstrong*, Boston: D. Reidel, pp. 191–223.

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

Earman, J. and Roberts, J. [2005b]: 'Contact with the Nomic: A Challenge for Deniers of Humean Supervenience about Laws of Nature (Part II)', *Philosophy and Phenomenological Research*, **71** (2), pp. 253–68.

Elga, A. [2004]: 'Infinitesimal Chances and the Laws of Nature', *Australasian Journal of Philosophy*, **82** (1), pp. 67–76.

Goodman, N. [1979]: *Fact, Fiction, and Forecast*, New York: Harvard University Press.

Hall, N. [2004]: 'Two Mistakes about Credence and Chance', *Australasian Journal of Philosophy*, **82** (1), pp. 93–111.

Hall, N. [MS]: 'Humean Reductionism about Laws of Nature', ¡philpapers.org/rec/HALHRA¿ accessed September 13, 2013.

Hicks, M. T. and Schaffer, J. [2016]: 'Derivative Properties in Fundamental Laws', *British Journal for the Philosophy of Science online preprint.*

Hoefer, C. [2007] 'The Third Way on Objective Probability: A Sceptic's Guide to Objective Chance', *Mind* **116** (463), pp. 549–596..

Huemer, M. [2009]: 'When is Parsimony a Virtue', *Philosophical Quarterly*, **59** (235), pp. 216–36.

Jeffries, H. [1973]: *Scientific Inference, 3rd Edition*, Cambridge: Cambridge University Press.

Joyce, J. [2007]: 'Epistemic Deference: The Case of Chance', *Proceedings of the Aristotelian Society*, Vol. CVII, Part 2.

Kitcher, P. [1989]: 'Explanatory Unification and the Causal Structure of the World', in Kitcher, P. and Salmon, W. (*eds*), *Scientific Explanation*, Minneapolis: University of Minnesota Press, pp. 410–505.

Lange, M. [2009]: *Laws and Lawmakers: Science, Metaphysics, and the Laws of Nature.* Oxford: Oxford University Press.

Lewis, D. [1980]: 'A Subjectivist's Guide to Objective Chance', in Jeffrey, R. C. (*ed*), *Studies in Probability and Inductive Logic*, Berkeley: University of California Press.

Lewis, D. [1983]: ‘New Work for a Theory of Universals’, in *Papers in Metaphysics and Epistemology*, Oxford: Oxford University Press.

Lewis, D. [1994]: ‘Humean Supervenience Debugged’, *Mind*, **103** (412), pp. 473–90.

Loewer, B. [1996]: ‘Humean Supervenience’, *Philosophical Topics*, **24** (1), pp. 101–27.

Loewer, B. [2007]: ‘Laws and Natural Properties’, *Philosophical Topics*, **35** (1/2), pp. 313–28.

Loewer, B. [2008]: ‘Why There Is Anything Except Physics’, in Hohwy, J. and Kallestrup, J. (*eds*), *Being Reduced: New Essays on Reduction, Explanation, and Causation*, Oxford: Oxford University Press.

Loewer, B. [2009]: ‘Why Is There Anything Except Physics?’ *Synthese*, **170** (2), pp. 217–33.

Loewer, B. [2012]: ‘Two Accounts of Laws and Time’, *Philosophical Studies*, **160** (1), pp. 115–37.

Maudlin, Tim [2007]: *The Metaphysics Within Physics*, Oxford: Oxford University Press.

McKenzie, K. [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.

Mill, J. S. [1843]: *A System of Logic, Ratiocinitive and Inductive*, Reprint, New York: Echo Library 2009.

Mitchell, S. [1997]: ‘Pragmatic Laws’, *Philosophy of Science*, **64**: S468–S79.

Mitchell, S. [2000]: ‘Dimensions of Scientific Law’, *Philosophy of Science* 67: 242–265.

Nolan, D. [1997]: ‘Quantitative Parsimony’, *British Journal for the Philosophy of Science*, **48** (3): 329–43.

Ramsey, F. P. [1927]: ‘Laws and Causality’, reprinted in D. H. Mellor (*ed*) *F. P. Ramsey: Philosophical Papers.*, Cambridge: Cambridge University Press 1990.

Roberts, J. T. [2009]: *The Law Governed Universe*, Oxford: Oxford University Press.

Tooley, M. [1977]: ‘The Nature of Laws’, *Canadian Journal of Philosophy*, **7** (4), 667–98.

Woodward, J. [2013]: 'Simplicity in the Best Systems Account of Laws of Nature', *British Journal for the Philosophy of Science*, **0**, pp. 1–33.

Woodward, J. [2013]: 'Laws, Causes, and Invariance', in S. Mumford and T. Matthew (*eds*) *Metaphysics and Science*, Oxford: Oxford University Press, pp. 48–72.