Best-System Laws, Explanation, and Unification[1]

Thomas Blanchard, University of Cologne

Preprint – forthcoming in *Humean Laws for Human Agents*, edited by M. Hicks, S. Jaag and C. Loew, Oxford University Press.

In recent years, an active research program has emerged that aims to develop a Humean best-system account (BSA) of laws of nature that improves on Lewis's canonical articulation of the view. Its guiding idea is that the laws are cognitive tools tailored to the specific needs and limitations of creatures like us. While current versions of this "pragmatic Humean" research program fare much better than Lewis's account along many dimensions, I will argue that they have trouble making sense of certain key features of the practice of fundamental physics. Indeed, these features seem to go against the very idea that laws are useful for agents like us. In my view, Humeans can address these issues by paying more attention to the explanatory role of laws. Following this idea, I will propose an account on which what makes a systematization the best is a kind of explanatory power, understood along the lines of the unificationist theory of explanation. The resulting view, I will argue, can make sense of those features of laws that other pragmatic accounts of laws have trouble explaining.

7.1

The guiding idea of the BSA is that the laws are the members of the systematization of the "Humean mosaic" that fares best with respect to certain theoretical standards – the "Humean mosaic" being the complete set of particular, non-modal matters of fact about the universe. The BSA is a version of Humean reductionism about laws of nature (the view that laws reduce to the Humean mosaic) as it posits no metaphysical structure over and above the Humean mosaic: the laws are nothing more than summaries of the mosaic that have certain desirable theoretical features.[2] A key question in the debates over the BSA is whether such a metaphysically lean view of laws can still make sense of their characteristic functions such as enabling induction, supporting counterfactuals, etc. Another important question – and my focus in this paper – is what makes a systematization "the best". In Lewis's (1983) canonical version of the BSA, the best system is the one that best balances strength (understood as the amount of information that the system by itself provides about the mosaic) and simplicity (understood as a syntactic property of the system). Today it is widely agreed is that this proposal is not quite right, and merely a first pass. Accordingly contemporary defenders of the BSA have proposed various criteria with which to replace

or supplement Lewis's standards (Albert, 2015; Cohen and Callender, 2009; Dorst, 2019b; Hicks, 2018; Jaag and Loew, 2020; Loewer, 1996, 2007). While they differ from one another in important ways, these proposals all converge on the idea that what makes a

[1] Thanks to Michael Hicks, Jenann Ismael, Siegfried Jaag, Christian Loew, and audiences at the January 2016 Workshop in Philosophy of Physics at the University of Arizona and the 33rd Conference in the History and Philosophy of Science at CU Boulder for extremely helpful comments and discussion.

[2] In this paper I am using 'Humeanism' and 'BSA' more or less interchangeably. There are in fact other versions of Humeanism than the BSA, but none as plausible or popular.

system “the best” is its usefulness for cognitively limited and practically oriented creatures like us.

For anti-reductionists, this pragmatic move may seem beyond the pale. Lewis himself was very concerned to avoid the charge that the BSA makes the laws relative to us.[3] But properly executed, a pragmatic take on the BSA need not yield any of the absurd consequences one may fear. (For instance, it need not entail that we can change the laws at will.) Moreover, from a Humean point of view, two considerations make it attractive to introduce pragmatic criteria into the BSA.

The first has to do with a particular challenge for the Humean – the challenge of explaining why the search for laws occupies such a central place in fundamental physics (see Hall 2012, 39-1). Anti-Humeans about laws can easily explain why physicists care so much about the laws: the laws are (or are grounded in) fundamental features of reality that govern how nature behaves, and this makes them automatically worthy of physicists' attention. But of course, on a Humean standpoint the laws are not part of fundamental reality or metaphysically privileged in any way. So Humeans must find some other explanation for our interest in the laws. And an obvious idea (in fact, perhaps the only plausible one available to Humeans) is that knowing the laws is pragmatically beneficial for agents like us. Humeans are therefore well-advised to endorse a pragmatic reading of the BSA, on pain of making it mysterious why the laws matter to us.

Second, a pragmatic approach is well-poised to address one of the main objections against Lewis's BSA. As shown by Hall (2012, 2015), Roberts (2008) and Woodward (2014), physicists do not value strength and simplicity as Lewis understands these notions, nor do they trade off strength and simplicity in the manner envisioned by Lewis. Thus Lewis's account doesn't match how physicists themselves think about the laws. This objection has a particular sting to it, as one of Lewis's major selling points for his account that made it attractive to many was its supposed fit with scientific practice (Lewis 1983, 41). But once we understand the laws as designed for agents like us, this mismatch between the BSA and scientific practice seems to largely disappear. Consider for instance the fact that physicists value fundamental theories that are compatible with many different possible initial conditions and are in that respect very uninformative. (Jaag and Loew (2020) call this feature of the laws “modal latitude”.) If the goal of the best system is to elegantly encode as much information about the mosaic as possible, as Lewis's BSA claims, this makes little sense. But suppose instead that the laws are designed in part to help us predict the behavior of the many subsystems of interest present in our physical environment. Given that such systems all start in very different states, it is no surprise that good candidates for the status of laws are expected to be compatible with a wide range of initial conditions. (Another reason why the laws must have modal latitude to facilitate predictions is that limited agents like us can rarely if ever know the exact initial conditions of a system (Jaag and Loew 2020, 16).) In a similar vein, defenders of the pragmatic approach to the BSA have shown that it can make sense of many other features of the laws that are not predicted by Lewis's account, e.g. the fact that laws can be tested independently of one another (Hicks 2018), or are expected to display certain symmetries (Dorst 2019b).

[3] Cf. his discussion of the ‘ratbag idealist’ (Lewis 1994, 479).

Clearly, then, a pragmatic approach to the BSA is the way to go for Humeans. But some hurdles remain to be cleared. While existing pragmatic versions of the BSA fit scientific practice considerably better than Lewis's account, they still fail to capture two key aspects of the way in which physicists evaluate candidate best systems. Those two features of fundamental physics are especially problematic for pragmatic Humeanism, as they appear to go against the very idea that the laws are tailored to be useful to limited agents like us. In fact, these features seem to make more sense on an anti-Humean view of laws, thus raising the worry that with respect to fitting scientific practice, it is anti-Humeanism that has the upper hand.

The first feature of fundamental physics that is problematic for pragmatic Humeanism is the fact that physicists aim for a "theory of everything" (TOE): a complete, all-encompassing theoretical framework that can in principle account for every physical phenomenon in the universe. This has been an especially salient and distinct feature of fundamental physics since Newton. The ideal of a TOE plays such an important role in fundamental physics that any theory that fails to account for a certain range of physical phenomena is automatically deemed non-fundamental, even if it is otherwise empirically successful. For instance, "effective field theories" that are highly predictively accurate at a certain level are regarded as non-fundamental because they break down at certain energy scales. Moreover, this completeness requirement appears to take precedence over other criteria for laws. For instance, simplicity only comes into play as a criterion of choice between theories that hold the promise of being able to account for all physical matters of fact whatsoever (with the possible exception of the initial conditions of the universe). As Woodward (2014, 102) notes, this is one upshot of Einstein's remark that "the supreme goal of all theory is to make the basic irreducible elements as simple and as few as possible *without having to surrender the adequate representation of a single datum of experience*" (1934, 165, my emphasis). A best-system account has better incorporate such a requirement, then, on pain of failing to match a key aspect of scientific practice. But pragmatic versions of the BSA currently on the market do not include any requirement of completeness.[4] Thus, the accounts of Dorst (2019), Hicks (2018) and Jaag and Loew (2020) – the most detailed pragmatic versions of the BSA currently available – all leave room for the possibility that the best system might fail to account for portions of the mosaic as long as it fares substantially better than its competitors in other useful aspects (e.g. by being more easily confirmable, or more computationally tractable, etc.). Yet physicists do not seem willing to consider this possibility.

It is no surprise that those accounts do not capture the premium that physics puts on completeness. From a pragmatic Humean standpoint, this aspect of the practice of physics is somewhat of an enigma.[5] To illustrate, suppose you are trying to predict the future behavior of a rock sitting on your desk.[6] Physicists expect the laws to be complete in the sense that, for any possible exact initial conditions of the rock, those laws can in principle predict the future behavior of the rock down to its minutest microphysical details.

[4] Nor did Lewis's account. Indeed, it couldn't: given the way Lewis understands strength, a comprehensive theory would have (absurdly) to provide a complete description of all that happens in the universe.

[5] But see Dorst (this volume) for a discussion of how pragmatic Humeanism could accommodate this feature of fundamental physics.

[6] An example of Elga (2007) also discussed by Jaag and Loew (2020).

(Likewise for every other physical subsystem, of course.) But to do this the laws need information about every point of an enormous spatiotemporal region, e.g. the entire cross-section of the rock's backward light-cone in relativity. This is much more information than we can ever hope to gather, and calculating the behavior of the rock on its basis would be unfeasible for us anyway. Compare with macro-generalizations like "massive objects on a stable surface stay at rest unless pushed" or "objects sitting unstably at the border of a table have a high probability of falling." These yield only coarse-grained predictions of the rock's macroscopic behavior and break down for some of the rock's possible initial conditions (e.g. thermodynamically abnormal ones). But at least they *can* actually be used by agents like us, and generally yield highly reliable predictions. More generally, when modeling the behavior of complex systems one can often achieve enormous gains in representational and computational tractability at a small cost in accuracy and predictive power (Dennett, 1991). If the search for laws is driven by the need to identify generalizations useful to agents like us, the premium put on completeness by fundamental physics is therefore mysterious, at least *prima facie*, since it comes at a considerable cost in tractability and user-friendliness.[7] (Pragmatic Humeanism here seems to make better sense of the practice of the special sciences.[8] Since the behavior of objects of those sciences are not feasibly representable or predictable in all of their microphysical details, the common strategy in those sciences is to sacrifice accuracy and scope by constructing generalizations and models that represent the behavior of those systems only roughly and approximately, but in a way that is tractable and usable by limited agents like us.)

Here is another feature of fundamental physics that is problematic for pragmatic Humeans. As John Roberts (2008: 16-24) observes, the notion of law at work in fundamental physics is highly selective: physicists sharply distinguish between fundamental laws and regularities that are "striking and pervasive" but nevertheless not a matter of fundamental law. Roberts's examples of the second category include the second law of thermodynamics, astronomical regularities such as Kepler's rules or the fact that all planets orbit the Sun in the same direction, and global cosmological facts such as the cosmic microwave background or the large-scale flatness of the universe. Roberts goes on to argue that because those regularities are highly informative, Lewis's BSA has trouble explaining why scientists do not regard them as laws. A similar problem besets existing pragmatic Humean accounts. While they recognize a variety of uses for the laws, these accounts all put the emphasis on the laws' ability to help limited agents like us make easy and reliable predictions on the basis of the limited portions of the mosaic we can observe. (On the pragmatic Humean approach, this focus on prediction makes sense, as creatures like us obviously have a clear practical interest in being able to make speedy and accurate

[7] Jaag and Loew (2020: 11-14), who recognize the issue, offer various considerations in response. For instance, they point out that the laws are 'error-tolerant' in the sense that small errors in the specification of initial conditions usually lead only to small errors in predictions (cf. Callender 2017: ch. 7). This makes the laws applicable even if we have only incomplete information about initial conditions. While this does alleviate the issue somewhat, the point remains that the laws could be made even more user-friendly with a small sacrifice in comprehensiveness, and that from a pragmatic Humean standpoint it is mysterious why physicists are unwilling to make such tradeoffs.
[8] Here I echo some remarks of Woodward (2014: 119), who argues that simplicity/strength tradeoffs posited in Lewis's BSA are more characteristic of the special sciences than of fundamental physics. See also Frisch (2014).

predictions.) And this makes it puzzling why scientists do not regard the regularities cited above as fundamental laws, given how predictively useful they are. Call this the *problem of selectivity.*

For the sake of illustrating the problem, assume that the laws of our world are those of classical mechanics, as was once believed to be the case. Note that adding (say) the second law of thermodynamics to the laws of classical mechanics would yield a system much more predictively useful for agents like us than classical mechanics alone. After all, with the second law of thermodynamics in hand, one can effortlessly and reliably predict the behavior of an enormous number of physical systems of interest to us, e.g. that the cup of coffee in my hand will reach room temperature within the next hour or so. True, that information can also be extracted from the laws of classical mechanics, together with information about the initial conditions of the cup and the room. But predicting the behavior of my cup of coffee on the basis of classical mechanics is far more difficult than predicting it on the basis of the second law in two respects. It requires much more information about initial conditions - namely, information about the exact initial microstate of the system formed by the cup and its environment, or at least about the probability distribution over the possible initial microstates of that system. And calculating the behavior of the cup based on that information and the laws of classical mechanics is more computationally difficult and involves more steps than extracting it from the second law. Another way to predict the cup's behavior via the laws of classical mechanics would be to first use them to derive the second law from initial conditions of the universe, and then apply the second law to predict the cup's thermodynamic behavior. But of course this also requires more information and would be more computationally challenging than predicting the cup's behavior from a system that has the second law built in it right from the start. From the standpoint of current pragmatic best-system accounts, it is therefore mysterious why physicists are not willing to regard the second law of thermodynamics as an additional fundamental law over and above those posited by classical mechanics, since such an addition would help limited agents like us make faster and easier predictions.[9] (A similar case could be made, I believe, with other generalizations that physicists do not regard as fundamental, like the law of free fall and other principles of terrestrial dynamics, or astronomical regularities such as Kepler's rules or the fact that all planets orbit the Sun in the same direction.)

Now, like Lewis's account, all current pragmatic versions of the BSA include simplicity on their list of desiderata for best systems. So an obvious suggestion is that the gain in predictive usefulness obtained by adding the relevant generalizations to the laws of classical mechanics is more than offset by the resulting decrease in the simplicity of the system, and this is why physicists do not regard those generalizations as fundamental laws. But that thought is hard to square with the way in which these accounts understand simplicity and its pragmatic benefits. In many of those accounts, the simplicity desideratum is intended to exclude systems that agents like us couldn't possibly comprehend or manipulate, such as systematizations that list all the facts about the mosaic, or Lewis's "predicate F" system (see e.g. Albert 2015: 23). But that motivation doesn't apply in our example. While adding (say) the second law of thermodynamics to classical mechanics

[9] Frisch (2014b) makes a similar point in the context of a discussion of Albert's and Loewer's pragmatic version of the BSA (Albert 2000, 2015; Loewer 2007, 2012).

yields a slight increase in complexity, the resulting system is certainly not representationally or computationally intractable for agents like us. A pragmatic preference for simpler systems could also be motivated by the fact that simpler systems require fewer cognitive resources to be stored (Jaag and Loew 2020: 11). But while the system made of the second law and the laws of classical mechanics requires slightly more storage in long-term memory than classical mechanics alone, this seems a small price to pay for the resulting gain in predictive usefulness. Finally, the simplicity requirement is also often motivated based on considerations of user-friendliness: simpler theories should be preferred because they are easier to comprehend or enable easier and faster computations (e.g. Dorst 2019b: 896-7). But this rationale *favors* the addition the second law (and other non-fundamental generalizations such as the law of free fall) to the laws of classical mechanics, since as noted above doing so greatly streamlines the inference of an enormous number of facts.

Another suggestion is that these generalizations do not count as laws because they fail to meet a pragmatic criterion of breadth. We have an obvious interest in identifying generalizations that apply to as many subsystems as possible, and accordingly virtually all pragmatic versions of the BSA include a desideratum of breadth in their recipes for lawfulness.[10] Perhaps, then, the second law of thermodynamics doesn't count as a law because there are many subsystems of the universe to which it doesn't apply. (This is also true – and even more so – of the generalizations of astronomy or terrestrial dynamics.) But this suggestion faces several objections. For one thing, the range of application of the second law – from gases and cups of water to galaxies and black holes – is still extraordinarily broad. So the proposal only works if the breadth standard that a generalization must meet to count as a law is set very high. On a pragmatic picture it is not clear why such a stringent standard would have become part of the practice of physics. Moreover, it is not entirely clear that lawful generalizations would meet that standard. The laws of classical mechanics *in principle* apply to all physical systems[11], but in practice their breadth is severely limited: there are plenty of physical systems that are too complex for agents like us to be able to predict their behavior on the basis of the laws. And on pragmatic Humeanism, presumably it is breadth in practice rather than in principle that would seem to matter. Finally, the proposal arguably leaves it mysterious why fundamental physics draws such a sharp and rigid distinction between regularities that are fundamental laws and regularities that are not. If lawfulness is so intimately tied to predictive breadth, it would make more sense for physicists to adopt a graded and context-sensitive notion of law, so that generalizations would count as more or less lawful depending on how broad they are, and/or which generalizations are laws would vary with the context. (For instance, Kepler's rules would count as laws in contexts where we are especially interested in astronomical predictions but not in others.)[12]

[10] For instance Dorst (2019b: 887-9), Hicks (2018: 997-8) and Jaag and Loew (2020: 19-20). This requirement helps those accounts explain why certain global facts about the universe such as its total number of particles do not count as laws despite being very informative in Lewis's sense.

[11] Remember that we are working on the assumption that they are the fundamental laws.

[12] Interestingly, such a flexible picture of laws is explicitly endorsed by at least some pragmatic Humeans, namely Cohen and Callender (2009). Though that picture seems to me to be significantly at odds with how physicists themselves think of laws, it is worth noting that it seems to fit the situation in the special sciences (as Cohen and Callender note in favor of their account). In the special sciences one finds predictive models that employ

A final response is that the second law *is* a fundamental law after all. Albert and Loewer argue, on pragmatic Humean accounts, that the low-entropy start of the universe (the "Past Hypothesis" or PH) and the statistical-mechanical probability distribution conditioned on the PH are fundamental laws (Albert 2000, 2015; Loewer 2007 2012). Together with the dynamical laws, these two facts entail the second law of thermodynamics (and arguably all other special science generalizations). This proposal makes the second law a theorem of the best system and thus a fundamental law or at least a direct consequence of the fundamental laws. But while one does find scattered remarks that fit with that view in the physics literature[13], this proposal doesn't seem to have taken much hold among physicists. Perhaps this is because physicists are still in the grip of an antiquated conception of laws that prevents them from recognizing the lawful status of PH and PROB. But those not yet firmly committed to Humeanism may not find this claim especially plausible.

To sum up: pragmatic Humeanism, in its current incarnations, has trouble making sense of the premium that fundamental physics puts on comprehensiveness, and of the fact that physicists recognize so few generalizations as laws. And the issue is an especially pressing one. For these two features of laws seem to go against the general idea that the laws are tailored to be useful to agents like us. But as noted above, on a Humean picture it is not clear how to make sense of physicists' interest for the laws other than through their supposed pragmatic benefits. So the problem here is one for Humeanism *tout court*. Moreover, anti-Humean views of laws (some of them at least) seem better poised to make sense of the features of laws under consideration. On a governing view of laws, for example, it makes sense to expect the laws to apply to every part and parcel of the universe. One can perhaps make sense of a scenario in which the laws govern only some portions of natural reality, though some complications would arise in developing that picture. (For example, what happens in interactions between those components of reality that are governed by laws and those that are not?) But in any case, that hypothesis seems more contrived and unnatural than the scenario on which the laws govern all of nature, so that it is no surprise that physics takes the latter as working hypothesis. A governing view can also easily explain why the second law of thermodynamics is not a fundamental law in addition to those of classical mechanics: the second law is an enormously and informative generalization, but it is the laws of classical mechanics that do the real work of governing thermodynamic phenomena. If right, this undercuts one the main considerations that Humeans routinely offer in favor of the BSA, namely that it fits scientific practice better than competitors. So the issue here is one that Humeans have to take seriously.

7.2

Fortunately, further reflection on the problem of selectivity suggests the beginning of the solution. I have argued that current pragmatic BSAs do not get the distinction between laws and non-laws right, and count too many generalizations as laws. But even if they did get the right results, it would be for the wrong reasons. Physicists do not decide

---

generalizations of varying degrees of breadth and invariance (see Woodward 2003), which play the functions characteristic of laws (explanation, counterfactual support, etc.) in some contexts but not others, and with little effort being made to rigidly separate those generalizations into 'fundamental' and 'non-fundamental' ones.

[13] For instance, Feynman (1965) envisions the possibility that the PH may be a law.

whether a generalization is a law based on how useful it is for limited agents like us to make predictions. Instead, I suggest, they rely primarily on *explanatory* considerations. Consider the fact that the second law of thermodynamics, which was generally regarded as a fundamental law in early 19th-century theories of heat, had lost that status among physicists by the end of the century (see Brush 1976). The reason was that by that time it had become clear that entropy increase could be explained as a consequence of the laws of classical mechanics. (Not, of course, as a consequence of those laws alone, but as a consequence of the laws together with facts about the initial state of the universe, as became clear through Boltzmann's work.) By contrast, the principles of classical mechanics *were* regarded as fundamental laws at that time because it was widely thought that, while being able to explain a great number of phenomena, they could not themselves be explained in terms of deeper principles. Further support for that conjecture comes from other examples of striking but non-lawful generalizations mentioned in the previous section. Quite evidently, the reason why physicists do not regard, say, Kepler's rules as fundamental is that they can be explained as consequences of classical mechanics and the initial conditions of the solar system (including among other things the absence of any body at least as massive as the sun in the vicinity of the solar system). Likewise, the uniformity in planetary orbital directions, the cosmic microwave background and the flatness of the universe can all be explained as consequences of generalizations we already regard as fundamental and facts about initial conditions (respectively the origins of the solar system as a gigantic gas of dust and cloud, the Big Bang, and the inflationary period).

These considerations suggest the following hypothesis. No matter how informative, broad or useful a generalization is in itself, physicists will not deem it fundamental if it can be explained through generalizations that they already regard as fundamental laws. And conversely, they will deem it fundamental if it is not itself explainable in terms of further facts. In other words, physicists take the fundamental laws to be the *ultimate explainers* – the principles that explain other facts about the universe, while not themselves being explainable in terms of deeper principles. Hence a proposal for Humeans: an adequate BSA should take as starting point the idea that the chief virtue that makes a system best is a kind of *ultimate explanatory power*.

As I explained, the key motivation for this proposal is that it seems to get the distinction between laws and other regularities right – at least more so than other Humean accounts. (It also fits tightly what some physicists have to say about laws. Weinberg, for instance, speaks of the quest for a final theory as "the ancient search for those principles that cannot be explained in terms of deeper principles" (1992: 18).) But questions and concerns immediately arise. Explanation is an elusive notion, and using it as starting point for a Humean account of laws raises a number of issues. Many Humean-friendly theories of explanation – most obviously the deductive-nomological account (Hempel 1965) – presuppose a distinction between laws and non-laws already in place and hence cannot be used to elucidate what laws are. Furthermore, it is unclear how Humeans can make sense of the laws as *ultimate* explainers, since Humeanism also holds that laws do have explanations – they are metaphysically explained by their instances.[14] Finally, even if the

[14] This issue is related to the well-known problem of "explanatory circularity" for Humeanism: if laws are just regularities, they are explained by their instances and hence cannot themselves figure in explanations of those instances, on pain of explanatory circularity (see e.g. Lange 2013). In my view the best answer to that

proposal fits with how physicists themselves think of laws, it is unclear whether it can do better job other Humean accounts at motivating puzzling aspects of scientific practice, such as the ones discussed in section I. On an explanationist best-system account, the search for a comprehensive theory is naturally understood as a search for a theory that can (in principle) explain everything, but what exactly is the value of such a theory? And granting that physicists' distinction between laws and other striking regularities is based on explanatory considerations, what exactly is the payoff of drawing such a distinction? The underlying issue is that whereas making good predictions has an immediate and tangible value for agents like us, the point or value of explanation is rather opaque, at least from a Humean perspective.[15] In Salmon's words: 'Why ask, "Why?"?' (Salmon, 1978). One answer is that explanations provide us with a sense of understanding, but this raises more questions than it answers. So however well it fits actual scientific practice, an explanationist take on the best-system approach runs the risk of making it mysterious why physicists care so much about the laws.

In the remainder of the paper I will sketch how an explanationist BSA can be developed to address these questions and concerns. In the next section I will propose a way to spell out the account in terms of explanatory unification. In section 7.4, I will argue that the proposal can make good sense of why physicists care so much about the laws, and can also make sense of the features of laws discussed above that other pragmatic Humean accounts of laws leave mysterious.

7.3

Several considerations suggest unification as a key notion to develop an explanationist best-system account. The idea that explanation involves unification has a long pedigree in the philosophy of science (where it is associated mainly with Friedman (1974) and Kitcher (1981, 1989). And there is something deeply intuitive to the idea. Unification involves drawing together seemingly unrelated under a single cohesive theoretical framework, and such an achievement seems explanatory valuable. In addition, even a cursory glance at the history of physics reveals a prominent role for considerations of unification. Many theories that physicists take or once took seriously as fundamental gained that status largely because they unified previously disparate domains: e.g. celestial and terrestrial dynamics for Newtonian mechanics, electricity and magnetism for Maxwell's equations and special relativity, and electromagnetism and weak nuclear theory for the electroweak theory. And contemporary physics' search for a TOE is fundamentally a search for unification - first, of the electroweak and strong forces into a Grand Unified Theory (GUT), which is then itself to be unified with gravity. This strongly suggests that the ultimate explanatory power that physicists expect from fundamental laws should be understood as unification. Finally, unificationism about explanation and the best-system account of laws go very well together. What makes a system best for the BSA is its ability to cover a wide range of phenomena via a suitably simple set of principles, and this is also

objection is that it runs together two distinct kinds of explanation, scientific and metaphysical (Loewer 2012; see also Dorst 2019a and Bhogal 2020). I will briefly return to this issue in the next section.

[15] Non-Humeans may claim that explanation has an obvious value: its function is to track metaphysically robust relationships of dependence in the world. But of course Humeans deny the existence of such relations.

what explanatory unification involves. Accordingly, several defenders of the BSA appeal to unification to make sense of the explanatory role of laws (e.g. Loewer (1996: 113) and Bhogal (2020)[16]), though they do not go as far as claiming that explanatory unification is what *makes* a system the best, as I am proposing here.[17]

Before explaining further how this goes, let me clarify that I do not wish to claim, as for instance Kitcher does, that *all* explanation is unification. Such an imperialist view of unification faces several objections. For instance, it entails that causal explanation can be reduced to unification. But it is hard to see how unification can account for the asymmetry of causal explanation (Barnes 1992). And Kitcher's attempt to fit causal explanation within a unificationist framework forces him to endorse an implausible 'winner-take-all' on which only the most unifying theories are explanatory at all (Woodward 2017). I want to be a pluralist about explanation, and claim only that the explanatory role *of fundamental physical laws* can be understood in terms of unification. Explanation in the special sciences, on the other hand, is mostly causal and does not aim at unification.[18] This modest position escapes the objections against unificationism just mentioned.

(In fact, I do not even wish to claim that the explanatory role of fundamental physical laws is exhausted by unification. Physicists presumably rely on a diverse range of considerations when assessing the explanatory value of a fundamental theory. A fully developed explanationist best-system account is thus likely to include additional criteria on good systematizations besides unification, though I will not explore what these might be here.)

My way of understanding unification is inspired by Friedman's (1974) account, though different from it in important respects. The main idea is as follows. One starts with a set S of facts to be unified. A theory T consists of a subset of those facts. We can think of the facts included in T as the basic theoretical principles or laws of the theory. The degree of unification of S by T is a function of how sparse T is (i.e., how many laws it contains), and of how many facts in S can be derived from T in a certain way. (I will explain what counts as a proper derivation shortly.) T is maximally unifying when it can properly derive all the facts in S, and no sparser theory exists that does the same. When a theory is maximally unifying, all facts in S can be derived from its laws, and those laws themselves cannot be derived from deeper, more encompassing principles. As should be clear, this fits closely with the idea of laws as ultimate explainers. My proposal, then, is that what makes a systematization the best is that it maximally unifies the mosaic, or at least comes closest to the ideal of maximal unification. (The reason for that caveat will appear later.)

To make this more precise, we need to specify the set S of phenomena that the laws are intended to unify. That set includes all the fundamental physical facts that compose the Humean mosaic. We need not take any specific stance on the nature of those fundamental

[16] Bhogal uses the idea that laws explain by unifying to distinguish scientific explanations from metaphysical ones (which do not aim at unification) and to thereby answer the explanatory circularity objection to Humeanism. This response fits very well with the account I propose in this paper.

[17] Psillos (2002: ch. 10) defends a similar view, though with motivations different from mine. Kitcher (1989) also seriously considers the possibility that lawfulness can be understood on the basis of a prior notion of explanatory unification, and offers an interpretation of Kant's theory of laws along these lines (Kitcher 1996).

[18] Of course, on my proposal considerations of unification comes into play in determining what the laws are and hence also in (partly) fixing the causal facts that work as *explanantia* in causal explanations. But this impact of unification on causal explanation does not entail that causal facts themselves do explanatory work by unifying.

physical facts here.[19] We must also include in *S* all the facts that supervene on the Humean mosaic. This is in line with the idea that physics aims to explain not only physical phenomena but also all other phenomena that depend on the physical. (Also, Humean laws are regularities that supervene on the mosaic. Since a fundamental theory is supposed to be a subset of the facts in *S* that set better include the laws.)

We also need to explain what it means for a fact *f* in *S* to be properly derivable from a theory *T*. Requiring *f* to follow from *T* alone would be too strong. True, some particular facts about the universe directly follow from a fundamental theory. For example, the fact that a particular magnetic field has a net flux equal to zero can be deduced directly from Maxwell's second equation. But generally the laws only entail a particular fact only given additional information - information that we would naturally describe as being about 'initial conditions'. Consider for instance a proposition p_2 describing the position of a planet at a given time. To derive p_2 from the laws of Newtonian mechanics, one needs a further proposition p_1 describing positions, velocities and masses of the Sun and planetary bodies at an earlier time. I propose that we understand proper derivation as follows:

> A fact *f* in *S* can be properly derived from *T* just in case there exist facts $f_1, \ldots, f_n$ in *S* distinct from *f* such that $f_1, \ldots, f_n$ and *T* together entail *f*.[20]

When I say $f_1, \ldots, f_n$ must be *distinct* from *f*, I mean that *f* is conceptually and metaphysically independent of $f_1, \ldots, f_n$ (Lewis 1986). Without this qualification, we would get the disastrous result that the best system is the empty one, since such a system would be maximally sparse and could 'account' for every fact in *S* by deriving it from itself.[21] Note that this conception of derivation includes no built-in constraint that *f* be derived from *T* and facts at other times. But as Hicks (2018) notes, it is a plausible conjecture that, given the way that the contents of the universe are *de facto* distributed in spacetime, any theory that aims to derive a large number of physical facts while remaining tractable by agents like us will have to take the form of a dynamical theory that derives facts at a time from facts at other times.

Finally, to fully flesh out the proposal we need to explain when a theory counts as sparser than another. This raises a number of issues. For example, compare a system T_1 consisting of Galileo's laws of terrestrial dynamics and Kepler's astronomical rules with the system T_2 obtained by conjoining those laws into a single statement.

---

[19] Lewis, of course, takes them to be instantiations of perfectly natural properties at spacetime points, but that characterization raises a number of issues (Loewer, 2007).

[20] This conception of derivation as a kind of entailment works only in the case of deterministic theories. There are natural ways to extend the account proposed here to encompass indeterministic theories. But I will not discuss them here.

[21] The distinctness requirement also helps address a related objection. If a system *T* contains a law *L*, it seems that the system *T'* obtained by eliminating *L* from *T* always counts as better. For *T'* is certainly sparser than *T*, and can derive *L* (and hence also all the facts that *L* helps derive) by listing all of *L*'s instances. But since *L*'s instances are not distinct from *L*, this doesn't count as a proper derivation. This fits well with the idea that scientific and metaphysical explanations are importantly different, as many Humeans have argued in response to the explanatory circularity objection (see fn. 13). While *T'*'s derivation of *L* from its instances might count as a proper metaphysical explanation of *L*, it doesn't count as a scientific explanation of *L*.

Obviously T_2 is not sparser than T_1 in any meaningful way, but explaining why is not easy.[22] Friedman's solution to this problem appeals to the idea that genuine identification works by reducing the total number of 'independently accepted' regularities, but that notion faces severe difficulties (Kitcher, 1976). A better suggestion is that the statement obtained by conjoining Galileo's and Kepler's laws cannot account for any phenomena not already covered by these laws, and hence does not genuinely unify them into a single law.[23] Compare with the genuine unification of terrestrial and celestial dynamics provided by Newton's law of gravitation: that law can account for all phenomena covered by Galilean and Keplerian laws, as well as other phenomena, such as meteorites and other objects that lie at the boundary of those laws (Douglas 2009: 456). (Interestingly, on that proposal replacing several laws by a single law yields a genuinely sparser theory only if it also increases in the number of phenomena derivable from the resulting theory.) These remarks certainly do not amount to a full account of sparseness, but giving such an account goes beyond the scope of the paper.

To explain the logic of my account, let us assume that the laws of our world are those of Newtonian mechanics. On my view, what makes them the laws is, first, that all the facts about the mosaic can be properly derived from Newtonian mechanics, generally together with further information about other parts of the mosaic, in the manner that p_2 can be derived from p_1 (and p_1 itself would be derivable from even earlier facts about the solar system, and so on). Second, there is no sparser set of principles which can also derive all the facts in *S*. If there was one, then Newton's laws (which are among the facts in S) could themselves be properly derived from this deeper, more encompassing set of principles. It is because they cannot be derived in this way that they are the ultimate explainers. (The relevant sense of 'ultimate', note, is compatible with laws being grounded in their instances and hence not *metaphysically* ultimate. Replacing a law with a list of its instances would not yield a more unifying system. For instance, a system that simply listed all instances of Newton's laws with a list of their instances would be able to derive both those laws and all of the facts that Newtonian mechanics can properly derive, but would not be in any reasonable sense sparser than Newtonian mechanics.)

I will close this section by mentioning one complication for my proposal and explaining how to address it. (Readers pressed for time may safely skip to section IV.) The complication is that if the universe had a beginning its very initial conditions may not be explainable in terms of the laws, so that no fundamental theory can be the best system in my sense. Call this *the problem of initial conditions*. (One wrinkle is that many of the theories that physicists take seriously are time-symmetric in the sense that they permit the derivation of past from future facts. Those theories therefore can 'properly derive' initial conditions of the universe – namely, from *later* time-slices of the universe. However, physicists do not seem to regard such later-from-earlier derivations as explanatory. Perhaps there is in fact no privileged direction of explanation in physics and this is mere prejudice in their part. Or (more likely in my view) it is evidence that physicists rely on additional considerations besides unification when assessing explanations. At any rate, the time-symmetry of physical theories does not offer a clear way out of the problem of initial conditions.)

[22] See Hempel's discussion of the "problem of conjunction" (1965: 273).
[23] Here I am indebted to Psillos (2002: 272), who makes what I take to be the same suggestion.

The proper response to the problem, in my view, is that one can still hope that nature is kind enough to let a clear winner emerge from the competition between systems. In particular, suppose there is a system *T* from which every fact in *S* can be properly derived except facts about initial conditions of the universe. *T*, then, is a theory of nearly everything. And suppose also that serious competitors to *T* (i.e. alternatives to *T* that can explain a reasonable amount of facts in *S*) cannot explain those initial conditions either, and either fail to account for some other facts in *S*, or can account for all remaining facts but in a less sparse way than *T*. *T* would then clearly come closest to the ideal of a maximally unifying theory, and thus deserves the title of best system. Modern physics, while willing to consider the possibility that initial conditions may not be explainable, does seem predicated on the hope that a theory with the features just described exists.

Initial conditions raise another issue for my account, however. Assume again that the laws of our universe are those of Newtonian mechanics, and let *p* be a proposition describing the initial conditions of the universe. Supposing that *p* cannot be properly derived from Newtonian mechanics, my proposal would seem to require making *p* a fundamental law as the system made of Newton's laws and *p* would account for more facts than Newtonian mechanics alone. This is at odds with how physicists think of initial conditions. A sensible response to this worry, however, is that *p* is likely to be too complex a proposition for us to ever be able to represent it, so that a system that includes it as a law would be cognitively intractable by agents like us. What if *p* describes just one of the small, localized facts that together compose the initial conditions of the universe? However, counting *p* as a law goes very much against the spirit of the unificationist approach, since this would do little to increase the unification of the mosaic. The resulting law would be an *ad hoc* addition designed to explain just one little fact rather than drawing together seemingly disparate facts under a common umbrella, as *bona fide* laws do. A unificationist best-system account can therefore address the issue by requiring laws to earn their keep: Adding a new law to a system is permissible only if that new law helps the system account for a substantial amount of facts that could not be derived from it otherwise. I leave further development of this idea for future work.

7.4

I have argued that a proper best-system account should take as point of departure the idea that laws play a kind of ultimate explanatory role, and sketched a way to spell out such an account in terms of unification. I now want to discuss how the resulting account can help pragmatic Humeans make sense of the two problematic features of fundamental physics discussed at the outset of this paper. These, remember, were the fact that physicists aim for a theory of everything, and count only a few select generalizations as laws. These two aspects of scientific practice are built into my account, on which the best system must be able to account for all or nearly all phenomena in the mosaic, and must do so by positing as few laws as possible. Of course, that my account recovers these aspects of scientific practice doesn't mean that it explains them. The question remains of why physicists spend so much time and effort finding a theory with these features, despite its lack of user-friendliness for predictive purposes.

To answer that question, it helps to return to an issue for my account already mentioned at the end of section 7.2. This is the fact that the point of explanation is rather opaque, so that an explanationist version of the BSA runs the risk of leaving it mysterious why physicists care about the laws. One answer to that worry is that explanations are valuable because they provide understanding. Friedman (1974), for instance, endorses such a view in the context of defending his unificationist account of explanation. According to him, we gain understanding by reducing the number of independent facts we have to posit as brute. Assuming that physics seeks understanding, it therefore makes sense that physicists seek to identify the sparsest set of principles from which everything else can in principle be derived. Still, the question arises of why understanding itself is valuable. What exactly would be lost if we stopped seeking understanding?

A more promising line of response, in my view, is that explanations are instrumentally valuable, and that seeking and identifying explanations has tangible benefits for agents like us. In ordinary cognition, explanatory reasoning has been shown to have wide-ranging effects on a variety of cognitive tasks such as discovery, confirmation and learning (Lombrozo 2011; 2016). My suggestion is that Humeans can provide a story of this kind to explain why laws with explanatory/unificatory power are valuable: searching for such a system is an efficient and fruitful way to explore the universe.[24] To see this, note that many striking facts about the ground of physics have been discovered via explanatory inference – that is, by noting that these facts together with the laws explain certain observations. Consider for example Laplace's discovery that the solar system originated in a gigantic gas cloud. Laplace arrived at this 'nebular hypothesis' by reasoning that together with Newton's laws, this hypothesis would explain the puzzling fact that in the solar system all planetary orbits lie on a single plane and follow the same direction (a fact that Newton himself took to be the result of the direct intervention of God).[25] Likewise, striking facts about our early universe such as the Big Bang, the inflationary period, or the low entropy start of the universe have been established on the ground that together with the laws these facts explain various cosmological regularities - respectively the cosmic microwave background, the large-scale flatness of the universe, and the second law of thermodynamics. So we have a pervasive pattern in physics where some important and interesting fact E about the mosaic on the ground that E together with the laws L enables the derivation and hence the explanation of a seemingly puzzling observation O. (In my jargon, this amounts to showing that given E one can 'properly derive' O from L.) This pattern suggests that by playing their explanatory role the laws serve an important function for us: they help us extract information from our observations. An explanationist version of the BSA can therefore vindicate the pragmatic Humean idea that what makes a system the best is its usefulness for agents like us. Limited agents like us face a particular epistemic predicament: we can observe only a very small portion of the universe, and need to figure what the rest is like on the basis of those very restricted observations. The laws help us solve that predicament by enabling the kind of explanatory inference just described. No wonder, then, that physicists are so interested in discovering the laws.

Now, it remains to be shown that this line of thought is still plausible when explanation is understood as unification, as I have proposed. That is, we need to show how

[24] Dorst (2019a) has a similar proposal about how Humeans should understand the value of scientific explanation, though he develops it in a different way than I do here (partly because his account of explanation is less tied to unification than mine).
[25] See Roberts (2008: 17-9).

a system that is unifying in my sense –a system that allows the derivation of everything in the mosaic from a set of principles as sparse as possible – favors the exploration of the mosaic in the way just described. If such a story can be told, we would finally have a way to justify the premium that physicists put on comprehensiveness and sparseness in pragmatic Humean terms.

I cannot to offer a fully fleshed-out story of this form in this paper, but here is a sketch of how it may go. Return to the example of Laplace's nebular hypothesis. Before Laplace's work, it was not yet known whether and how the uniformity in planets' orbital directions (call that fact *U*) could be made to fit with the laws of Newtonian mechanics. Faced with this question, an 18th-century physicist could have adopted one of three strategies. The first would be to try and provide a Newtonian explanation of *U* by figuring out what kind of initial conditions could have led to this striking uniformity via Newton's laws. This was of course Laplace's strategy. A second strategy would be to posit that *U* does not admit of a natural explanation in terms of physical laws, thereby 'reconciling' *U* with Newtonian mechanics by placing the former outside of the latter's purview. This was in effect Newton's move. A third strategy would be to explain *U* by positing it as a new basic law in addition to those of Newtonian mechanics: if it is a fundamental law that all planets orbit the sun in the same direction, that immediately entails that they all do, and there is no further explanation to be sought. (Pragmatic versions of the BSA that emphasize the predictive role of laws validate that third strategy, since adding *U* to the book of laws would yield a substantial increase in predictive power.[26])

But now note that, if one wants to hold on to Newtonian mechanics as fundamental theory, the requirements of comprehensiveness and sparseness in effect privilege the first strategy. The rule that every physical phenomenon be properly derivable from the laws prohibits the second strategy. And the requirement that the best system be as sparse as possible rules against the third strategy, by prohibiting positing new laws lightly. It makes postulating a new law a solution of last resort, to be avoided unless there really is no hope of accounting for that phenomenon on the basis of principles already recognized as laws.[27]

Note, moreover, that it is only by pursuing the first strategy that one could be led to discover the facts about the origins of the solar system that explain *U*. (Indeed, our best reason to believe that the solar system started out as a rotating gas cloud is still that this together with the laws explains uniformities in planetary orbits.) Had the scientific community in the 18th century not pursued that strategy, and be content to either restrict the explanatory scope of Newtonian mechanics or posit *U* as a new fundamental law, that fact may well have never been discovered, or only much later.[28] Other examples with the same moral are easy to find. Think of other striking physical regularities that physicists do

---

[26] Perhaps *U* is not sufficiently predictively useful to deserve inclusion. But the point would still hold if we were to replace *U* with the second law of thermodynamics, or another extremely broad but still non-fundamental generalization.

[27] I note in passing that this proposal makes sense of the fact that physicists have yet to endorse Albert and Loewer's proposal that the Past Hypothesis is a law: the jury is still out on whether the PH can be explained in terms of the dynamical laws and other facts about the early universe.

[28] Roberts (2008: 21-2) argues that Lewis's BSA warrants positing *U* as a law and hence cannot explain why Laplace devoted so much effort to providing an explanation of *U*. This is part of his broader argument to the effect that Humean accounts of laws cannot make sense of physicists' distinction between laws and other regularities (cf. section I). In effect the account offered in this paper provides a response to this argument.

not regard as a matter of fundamental law – e.g. the red shifting of galaxies, the large-scale flatness of the observable universe, or the constant increase in entropy. In each case, attempts to explain those regularities on the basis of current fundamental physical theories led, via inference to the best explanation, to the discovery of further striking facts about the early history of the universe (the Big Bang, the inflationary period, and the Past Hypothesis). Again, had scientists been content to restrict the scope of their fundamental theories or posit those regularities as laws, they would likely have missed out on those discoveries.

Putting these considerations together suggests a hypothesis about the value of the comprehensiveness and sparseness requirements on the best system. Those requirements act as norms on the kinds of moves that physicists can make when it is unclear how a preferred fundamental theory *T* can accommodate some observation *O*. In effect, they encourage one possible reaction to that situation – namely, searching for further facts through which *T* can explain *O* – at the expense of other possible reactions. And the history of physics shows that in our world searching for such explanations is fruitful, and constitutes an excellent way to discover significant facts about the universe. It is therefore no surprise that the practice of fundamental physics, which aims at the discovery of such facts, is structured by such norms.[29]

Here is an objection against this line of argument. It is clear that investigation of the world must obey norms requiring scientific theories to be reasonably comprehensive and simple. If it were acceptable to restrict the scope of one's theory (or postulate a new principle *ad hoc*) whenever one encounters some recalcitrant phenomenon, scientific theorizing would likely not go very far, and certainly wouldn't lead to the kind of significant discovery exemplified by Laplace's nebular hypothesis. But why insist for *maximal* comprehensiveness and sparseness? A worry similar to the one raised in section I for other pragmatic Humean accounts arises here. After all, maximally comprehensive and sparse theories of the sort one finds in fundamental physics can *in practice* be applied only in a very few select contexts, so that there are many significant facts about the universe that we would want to know but could never hope to discover through those theories. For example, it would be ludicrous to hope to derive any interesting fact about, say, the actual evolutionary history of life on earth by doing fundamental physics.

But that objection can be put to rest if we pay attention to the fact that fundamental physics isn't all of science, and that the special sciences deploy a quite different strategy to investigate the world. More precisely, while the special sciences also display a concern for unification, each of them seeks to unify only one specific domain of reality (e.g. life, the mind, or the economy). And scientific theorizing in these disciplines is also heavily by considerations of user-friendliness, tractability, and practical relevance to agents like us. These further considerations lead to tradeoffs with considerations of unification. For instance, there is no expectation that a good special scientific theory be able to account for *all* phenomena even within its domain (because a fully comprehensive theory would likely

[29] A similar explanation can be given of the requirement that a scientific theory can properly explain a fact by deriving it from other facts only if the latter are distinct from the former. (This requirement, remember, is built into my account of unification.) One could "explain" *U* by deriving it from a list of its instances, but such a "metaphysical explanation" would not yield any interesting scientific discovery of the kind Laplace made. The distinctness requirement in effect prohibits cheap explanations of that sort.

be too complex to help us efficiently navigate the world). And likewise there is generally little pressure to make the theory as sparse as possible (since often a demand for sparseness interferes with considerations of user-friendliness and tractability). By contrast, the strategy adopted by fundamental physics to investigate the world involves pushing the concern for unification to its maximum, while other considerations take a backseat.

By adopting these different strategies, physics and the special sciences nicely complement each other. By yielding theories limited in range but easily tractable and usable by agents like us, the special sciences help us discover many significant facts about the Humean mosaic that we could never uncover by doing fundamental physics. For instance, we can uncover the most important milestones in the history of life by using the generalizations of evolutionary biology along with those of anatomy, molecular biology, and other areas of biology. But there are also significant facts about the universe that in all likelihood, limited agents like us could discover only by adopting the investigative strategy characteristic of fundamental physics.

One vivid illustration is the discovery of the Higgs boson. Its existence was first postulated in the context of the Weinberg-Salam Georgi electroweak theory. That theory seeks to provide a unified account of an enormous number of physical interactions (including at energy scales that we never encounter) in terms of a few basic principles, and thus nicely exemplifies physics' obsessive pursuit of unification. The theory entails that at very high energy regimes the W and Z bosons (which carry the weak interaction) and the photon (which carries the electromagnetic interaction) are all massless. This raises the question why at lower energy regimes the W and Z bosons become massive while photons remain massless. Postulating the Higgs boson provides an explanation of that fact by supplying a mechanism that breaks the symmetry between photons and W/Z bosons as temperatures decrease. This is a dramatic example of a significant discovery falling out of physics' strategy of searching for a maximally unified theory of the universe, and it is hard to imagine how we could have made that discovery other than by following this strategy.

In short: what justifies the norms of maximal comprehensiveness and sparseness is not that any fruitful inquiry into the natural world must abide by them. Instead, it is the fact that an investigative strategy guided by these norms enables the discovery of phenomena that no other investigative strategy available to us could uncover.

I will close by noting an interesting consequence of the argument presented in this section for the way in which Humeans should think of the value of laws. Most versions of the BSA assume that physicists seek to know the laws because the end product of that search is valuable. That is, we want to discover the laws because they contain information that is extremely useful to successfully navigate the world around us. In this picture the benefits of the laws kick in only once we know them. But consider Laplace's discovery of the nebular hypothesis. It would be wrong to say that Newtonian mechanics was known to be the true fundamental theory. For one thing, it isn't. And the question whether it could account for all phenomena in nature was still very much in dispute. In that context, the discovery of the nebular hypothesis served as a major piece of confirmation for Newtonian mechanics by showing that the theory could account for puzzling phenomena such as *R* without resorting to divine intervention. So the example holds an important lesson. On the view proposed in this paper, searching for the laws is a valuable scientific aim not only because the end product of that search is valuable, but also because the very *process* of searching for a maximally unifying account for the world leads to the discovery of significant facts about the mosaic. In other words, in our quest for the laws, it is not just the destination, but also the journey that matters.

References

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

Albert, D. (2015). *After Physics*. Cambridge MA: Harvard University Press.

Barnes, E. (1992). 'Explanatory Unification and the Problem of Asymmetry', *Philosophy of Science*, 59, pp. 558-71.

Bhogal, H. (2020). 'Nomothetic Explanation and Humeanism about Laws of Nature' in K. Bennett and D. Zimmerman (eds.) *Oxford Studies in Metaphysics,* 12. Oxford: Oxford University Press, pp. 164-202.

Brush, S. (1976). *The Kind of Motion We Call Heat*' Amsterdam: North-Holland Publishing Company.

Callender, C. (2017). *What Makes Time Special?* Oxford: Oxford University Press.

Cohen, J. & Callender, C. (2009). 'A Better Best System Account of Lawhood', *Philosophical Studies* 145, pp. 1-54.

Dennett, D. (1991). 'Real Patterns', *Journal of Philosophy* 88, pp. 27-51.

Dorst, C. (2019a). 'Humean Laws, Explanatory Circularity, and the Aims of Scientific Explanation', *Philosophical Studies* 176, pp. 2657-79.

Dorst, C. (2019b). 'Towards a Best Predictive System Account of Laws of Nature', *British Journal for the Philosophy of Science* 70, pp. 877-900.

Douglas, H. (2009). 'Reintroducing Prediction to Explanation', *Philosophy of Science* 76, pp. 444-63.

Einstein, A. (1934). 'On the Method of Theoretical Physics', *Philosophy of Science* 44, pp. 248-68.

Elga, A. (2007). 'Isolation and folk physics', in H. Price & R. Corry (Eds.), *Causation, physics and the constitution of reality*, Oxford: Oxford University Press, pp. 106–19.

Feynman, R. (1965). *The Character of Physical Law*. New York: Modern Library.

Friedman, M. (1974). 'Explanation and Scientific Understanding', *Journal of Philosophy* 71, pp. 5-19.

Frisch, M. (2014a). *Causal Reasoning in Physics*, Cambridge: Cambridge University Press.

Frisch, M. (2014b). 'Why Physics Can't Explain Everything', in Wilson, A. (Ed.), *Chance and Temporal Asymmetry*, Oxford: Oxford University Press, pp. 221-40.

Hall, N. (2012). 'Humean Reductionism about Laws of Nature', Unpublished manuscript. https://philarchive.org/archive/HALHRAv1

Hall, N. (2015). 'Humean Reductionism about Laws of Nature', in Loewer, B. and Schaffer, J. (Eds.). *A Companion to David Lewis*, Oxford: Blackwell, pp. 262-77.

Hempel, C. (1965). *Aspects of Scientific Explanation*. New York: Free Press.

Hicks, M. (2018). 'Dynamic Humeanism', *British Journal for the Philosophy of Science* 69, pp. 983-1007.

Jaag, S., & Loew, C. (2020). 'Making Best Systems *Best for Us*', *Synthese* 197, pp. 2525-50.

Kitcher, P. (1976). 'Explanation, Conjunction and Unification', *Journal of Philosophy* 73, pp. 207–12.

Kitcher, P. (1981). 'Explanatory Unification', *Philosophy of Science* 48, pp. 507-31.

Kitcher, P. (1989). 'Explanatory Unification and the Causal Structure of the World', in Salmon, W. and Kitcher, P. (Eds.), *Minnesota Studies in the Philosophy of Science*, vol. 13. Minneapolis: University of Minnesota Press, pp. 410–505.

Kitcher, P. (1994). 'The Unity of Science and the Unity of Nature', in Perrini, P. (Ed.), *Kant and Contemporary Epistemology*. Dordrecht: Kluwer, pp. 253-72.

Lange, M. (2013). 'Grounding, Scientific Explanation, and Humean Laws', *Philosophical Studies* 165, pp. 255-61.

Lewis, D. (1983). 'New Work for a Theory of Universals', *Australasian Journal of Philosophy* 61, pp. 343-77.

Lewis, D. (1986). 'Events', in *Philosophical Papers*, Volume II. Oxford: Oxford University Press, pp. 241-69.

Lewis, D. (1994). 'Humean supervenience debugged', *Mind* 103, pp. 473–90.

Loewer, B. (1996). 'Humean Supervenience', *Philosophical Topics* 24, pp. 101-27.

Loewer, B. (2007). 'Laws and Natural Properties', *Philosophical Topics* 35, pp. 313-28.

Loewer, B. (2012). 'Two Accounts of Laws and Time', *Philosophical Studies* 160, pp. 115-37.

Lombrozo, T. (2011). 'The Instrumental Value of Explanations', *Philosophy Compass* 6/8, pp. 539-51.

Lombrozo, T. (2016). 'Explanatory Preferences Shape Learning and Inference', *Trends in Cognitive Sciences* 20, pp. 748-59.

Roberts, J. (2008). *The Law-Governed Universe*. New York: Oxford University Press.

Salmon, W. (1978). 'Why Ask, "Why?"? An Inquiry Concerning Scientific Explanation', *Proceedings and Addresses of the American Philosophical Association* 51: 683–705.

Weinberg, S. (1992). *Dreams of a Final Theory*. New York: Pantheon Books.

Woodward, J. (2003). *Making Things Happen*. Oxford: Oxford University Press.

Woodward, J. (2014). 'Simplicity in the Best Systems Account of Laws of Nature', *British Journal for the Philosophy of Science* 65, pp. 91-123.

Woodward, J. (2017). 'Scientific Explanation', in Zalta, E. (ed.), *Stanford Encyclopedia of Philosophy*, Fall 2017 Edition.