The Package Deal Account of Naturalness

Harjit Bhogal

1 Introduction

In contemporary metaphysics the distinction between *natural* properties, like *mass* and *charge* and *unnatural* properties, like *grue* is very familiar (though, of course, not uncontroversial). The core idea, stemming from David Lewis's 'New Work for a Theory of Universals', is that this distinction is tightly connected to a variety of other important notions. Lewis identifies this distinction, and motivates the need to accept it, by arguing that we need to distinguish natural from unnatural properties in order to give adequate accounts of intrinsicality, laws, induction, causation, explanation, reference, and other notions. Work following this has identified even more roles for naturalness to play (e.g. Sider [2011], Dorr and Hawthorne [2013]). Further, Lewis, and much of the literature following, claimed that the distinction between natural and unnatural properties is primitive.

Again, since this is familiar ground, this is not the place to go into detail. What's important for our purposes is the way in which the distinction between natural and unnatural properties reflects the distinction between properties that legitimate to use in scientific theorizing and those that are not. (Let's, for now, restrict to fundamental science. That is, to fundamental physics, since the literature on natural properties has developed in a way that closely connects natural properties to those of fundamental physics.) It seems inappropriate for our basic physical theories to be built using extremely unnatural, gerrymandered properties, and this is because of the way that such unnatural properties shouldn't be used in inductions, in scientific explanations, in laws of nature and so on.

Here is something else that is familiar to metaphysicians, particularly to those reading this volume —

Lewis also developed a *Humean* approach to laws of nature. Humeanism in metaphysics is generally understood as the denial of necessary connections between distinct existences. A ball being scarlet can necessitate it being red, in a way that is consistent with Humeanism, because scarletness and redness are not distinct existences. But there cannot be such necessary connections between entities that are distinct. Of course, laws of nature are a challenge for such a view. If it's a law that F=MA then there seems to be a necessary connection between the distinct existences *force*, *mass* and *acceleration*. Lewis's view is one which can accept the existence of such laws, but deny that they imply any necessary connections.

The way this works is that we identify a fundamental base of the world that does not contain necessary connections and then we give an account of how the laws reduce to this base.[1] Lewis calls this necessary-connection-free base the *Humean mosaic*. The Humean mosaic consists in the intrinsic physical state of each spacetime point (or each pointlike object) and the spatio-temporal relations between those points. That is, the Humean mosaic is all the local property instantiations pushed up against each other in spacetime.

The laws are determined by this mosaic via the Best System Account (BSA). The basic idea of the BSA is that the laws are propositions that are relatively simple, but also informative about the mosaic. More precisely, consider sets of axioms. Some sets of axioms are informative about the mosaic – their deductive closure tells us a lot about the mosaics. Some sets of axioms are simple, in the sense of being syntactically simple when written down. The laws are the set of axioms that best balance simplicity and informativeness.[2] (Strictly, on Lewis's approach the laws are a subset of those axioms — the laws are the axioms which are universal generalizations. We will come back to this point later.)

Humean approaches of this kind are, as can be seen by this volume, very popular. Since Lewis, there has been a vast amount of work developing such approaches. Much of this work has been driven

[1] Strictly, Lewis doesn't use the notion of *reduction*, rather his claim is that the laws *supervene* on the mosaic. Though modern Humeans typically do talk about the laws reducing to, or being grounded in, or holding in virtue of, the Humean mosaic.

[2] Another dimension of goodness of a system – 'fit' becomes relevant if we are considering probabilistic laws. We will simplify things by only focusing on the non-probabilistic case in this paper.

by the ideas of Barry Loewer — via his own published papers, but also via his huge influence on a generation of philosophers working in this area.

One particularly interesting suggestion of Loewer's is that we can combine our account of laws with our account of natural properties. The BSA claims that a best system procedure can output the laws — the laws are axioms that best systematize the mosaic. Loewer suggests that natural properties can be outputted by a similar systemization procedure. In fact, the laws and natural properties arise from the same procedure. Roughly speaking, the natural properties are the properties that are referred to in the axioms that best systematize the mosaic. Consequently, Loewer calls this the *Package Deal Account*.

Just as the BSA was a huge advance in the metaphysics of laws — whether or not you like the BSA it did suggest a fully reductive account of laws which seemed to, broadly speaking, get the right results for what count as a law or not — the PDA, if successful, would be a huge advance in the metaphysics of naturalness. As we noted, the standard account of naturalness is a primitivist one — there is a basic distinction between natural and unnatural properties. A fully reductive account of naturalness would be a very substantial philosophical achievement.

However, the PDA is somewhat underexplored in the literature. Loewer's discussions of it in his [1996] and [2007] are only brief — although he has now developed the account in a bit more detail in his [fort.]. And there isn't much other written work on the account. (Eddon and Meacham [2015] are an exception – they survey a series of variants of the BSA which don't commit to a primitive notion of naturalness, including the PDA.) Further, Loewer's developments of the PDA differ in some substantial ways from each other — his thinking seems to have developed between 1996 and today.

But, even though the account is underexplored, it's still very influential. It's influence spreading not mainly via writing but rather, in the oral tradition, via informal discussion in the relevant communities of philosophers. Judging by the conversations that I've had, there is a large number of modern Humeans who think, or at least hope, that something like the PDA is the correct approach.

Given all this — the way the account is potentially important but underexplored, and the way that the ideas are influential but it isn't particularly clear how they are best developed — the PDA needs some exploration. And so in this paper I'm going to start this exploration. In particular, I'll map the connections between naturalness and laws on the traditional BSA and I'll consider how those connections have to adapted in order to develop different versions of the PDA. And I'll discuss how some versions of the PDA, including, perhaps, those favored by Loewer, are an instance of a larger approach to Humeanism — one that has been very visible in the recent literature — that focuses on the role of 'ideal observers' or 'ideal scientists'.

Since this paper is an exploration, my conclusion won't be very conclusive. But I'll suggest that these 'ideal scientists' approaches to the PDA are not successful, and while some versions of the PDA might be feasible they don't get you quite as much as you might have started out hoping for.

2 The Role of Naturalness in the BSA

But let's start by going back to the traditional BSA approach to laws. The way I'll describe the approach won't be very strictly Lewisian — the aim here isn't to accurately represent the nuances of Lewis's views, but rather to get on the table a pretty standard version of the BSA so that we can see the ways in which the PDA must deviate.

Again, the basic idea of Humeanism is that the world consists in the mosaic of events distributed across spacetime. And the BSA says that the laws are propositions that best balance simplicity and informativeness about this mosaic. In developing this account we can identify three distinct roles for naturalness to play.

(1) Naturalness plays a metaphysical role in characterizing the Humean mosaic. As we noted, the mosaic consists in the intrinsic physical state of each spacetime point (or each pointlike object) and the spatio-temporal relations between those points. For Lewis, though, intrinsicality is very closely connected to naturalness — so closely that the Humean mosaic is, in effect, the distribution of

natural properties instantiated at each space time point, and the spatio-temporal relations between these points. On the standard BSA the laws are reduced to this distribution of natural properties.

More generally, many people accept a connection between naturalness and *fundamentality* — the natural properties are fundamental, they are the basic building blocks of the world (see Tahko [2018, section 1] and Bennett [2017, section 5.7] for discussion).

So, the first role of natural properties in this picture is to help characterize the fundamental base.

(2) There is a role for naturalness to play in identifying what the relevant data is that theories need to be informative about. On the traditional BSA, a set of axioms scores better on informativeness when the axioms are informative *about the mosaic* — that is, about the distribution of natural properties across spacetime.

One way to put this is that there is a set of data sentences — it's informativeness about those data sentences that counts toward the bestness of systems — and on the traditional BSA those data sentences are restricted to sentences that say that certain natural properties are instantiated at spacetime points. So the data sentences are things like 'At spacetime point a properties x, y, z…are instantiated. (This, as we will see, is somewhat different from the way Lewis presented things, but it's useful to have the idea that naturalness can play a role in fixing the language of the data sentences in mind.)

There's an interesting question here about why the relevant data sentences are just about the mosaic. The basic intuition of the BSA is that the laws are those axioms that best balance simplicity and informativeness. So why would you restrict the data to the facts about the mosaic? Well, if you think that the facts about the mosaic determine all the other facts about the world, then there is at least one natural sense of informativeness where being informative about the world just consists in being informative about the mosaic. Of course, though, there are other reasonable senses of what it is to be informative about the world where a set of axioms being informative about the fundamental base does not imply that they are informative about the whole world. But for now, we can see that the metaphysical role of naturalness in characterizing the fundamental base of the world, motivates a certain conception of what the data is.

(3) The most discussed role of naturalness in the BSA is in fixing the language that the axioms can be formulated in. Part of what makes axioms better is that they are *simple*. And the traditional approach to measuring simplicity of axioms is *syntactic* — ideally there should be few axioms and they should be short. But to measure simplicity in this way we need to fix on a particular language — the length of an axiom depends upon the language it's expressed in, there is no reason to think that an axiom that is syntactically simple in English will be similarly syntactically simple in Urdu. So we need to identify some privileged language.

The first step, then, is to move to a formal logical language rather than a natural language, say the language of first-order logic. But, obviously, we need to augment the logical language with non-logical vocabulary — if the language has no non-logical vocabulary then we won't be able to formulate any informative statements about the mosaic.

But, here is the problem: Consider a predicate, F, that is instantiated by all and only the things that exist in the actual world (including, for example, spacetime regions and mereological fusions). Then a system with only one axiom, '$\forall x F(x)$', would be extremely simple. And also it would be, at least in one sense of informative, maximally informative because the truth of $\forall x F(x)$ rules out all non-actual worlds. But, clearly $\forall x F(x)$ should not count as a law. This is the *Predicate F Problem*. (Whether such systems really are informative is a question that we will consider more later.)

Lewis's solution is to restrict the language that the systems can be formulated in to a language where 'the primitive vocabulary…refer only to the perfectly natural properties' [Lewis, 1983b, p. 42]. This rules out $\forall x F(x)$ because the the predicate F does not refer to a perfectly natural property.

So this is another place naturalness comes into the traditional BSA approach — it help us fix the language that axioms are formulated in. And restricting the language in this way – to first order logic and predicates that refer to perfectly natural properties – seems somewhat fitting given the picture of the metaphysics as just natural properties being instantiated at spacetime points.

So, again, there are three 'slots' in this approach for naturalness to play a role, (1) in characterizing the fundamental base of the world, (2) in fixing the relevant data sentences — both what the data is about and the language the sentences are formulated in and (3) in fixing the language of the axioms.

Some Humeans, though, are unhappy with the central role that naturalness plays in this story about laws. In particular, many (but not all) Humeans are drawn to the view because of general empiricist or anti-metaphysical tendencies — for example, some Humeans want to avoid postulating necessary connections because they worry that such connections would be empirically indetectable (e.g. Earman and Roberts [2005]), or perhaps because such necessary connections would be strange or 'spooky' (what Maudlin [2007, p.71] calls 'prejudicial' motivations for Humeanism). And this suspicion towards postulating additional structure to the world carries over to natural properties. Many Humeans want to do without this appeal to a primitive distinction between natural and unnatural properties (see, for example, the views surveyed in Eddon and Meacham [2015]).

In particular, Loewer [fort.], following Demarest [2017], and van Fraassen [1989, p. 53] argues that the role of naturalness in the traditional BSA leads to the *mismatch problem*. Imagine that scientists have formulated something that they take to be a 'theory of everything'. The theory is widely accepted in the scientific community as the ultimate theory of the fundamental nature of the world. And the laws of such a theory very simply and informatively systematize the mosaic. According to the traditional BSA, though, it is possible that the laws of this theory of everything are not the actual laws of nature, because they are not formulated using perfectly natural properties. That is to say, the traditional BSA allows the possibility of there being a mismatch between the properties that are used in our best scientific practice and the natural properties. And the possibility of this mismatch might seem uncomfortable.

One way to put the concern is that the mismatch problem shows that the traditional BSA fails to show adequate respect for science. For example, Loewer [fort.] says that 'it seems presumptuous for a metaphysician to say to a physicist who believes she has found a theory that optimally satisfies all the scientific criteria but not the metaphysical one that she may not have discovered the laws since the theory is not formulated in the language of perfectly natural properties.' Another way to put the concern is that the potential mismatch between science and the natural properties casts doubt upon the empirical detectability of natural properties.

Either way, though, the mismatch problem points towards some of the same doubts that Humeans have with the existence of necessary connections. Defenders of the BSA often claim that their view respects science more than anti-Humean alternatives which postulate strange entities (e.g. Cohen and Callender [2009], [cite more]) and that there are issues with the detectability of necessary connections. The mismatch problem seems to be in the same spirit.

Because of these issues (and others that I'm not going to go into here) there is desire among some Humeans to build a version of the BSA that doesn't appeal to a primitive distinction between natural and unnatural properties (again Eddon and Meacham [2015] provide an excellent survey of such views).

But still, there does seem to be a distinction between natural and unnatural properties — there are some properties that are legitimate to use in scientific theorizing — that play the right roles in laws, causation, induction, explanation etc. — and some properties that are not. So, instead of just rejecting this, it would be ideal if this adapted version of the BSA can give an account of the distinction between natural and unnatural properties as well as giving an account of laws that doesn't take such a distinction as primitive. As we noted, that is the aim of the PDA.

Of course, though, in order for the PDA to succeed it needs excise naturalness from the traditional BSA. We will now look at each of the roles of naturalness we just identified in turn to see how, and if, the PDA can avoid the need for naturalness to play those roles.

3 Developing the PDA

3.1 Role (3) and the Predicate F problem

How can we develop the PDA? The idea, at it's simplest, is that we adapt the BSA so that the language of the axioms is not restricted to a set of pre-determined natural properties. Instead, we define the natural properties as the ones that are part of the best set of axioms — the ones that best balance simplicity and informativeness — and, further, those best axioms are the laws. That is to say, the PDA denies that natural properties play the role (3) that we identified earlier.

But now the Predicate F Problem rearises: If there are no restrictions on the language that the axioms can be formulated in then it seems like a system containing the single axiom $\forall x F(x)$ would count as the best system. Consequently, $\forall x F(x)$ would be the single law of nature, and F would be the single natural property — clearly not the desired result.

Loewer, however, doesn't think that we need to restrict the language of the axioms in order to rule out predicates like F. Rather, he thinks, $\forall x F(x)$ does not count as an axiom of the best system because it is *uninformative*, even though it's truth is inconsistent with all non-actual worlds: 'While [$\forall x F(x)$] might be maximally informative given Lewis's characterization of information as excluding alternatives this merely shows that Lewis's proposal for evaluating informativeness is not relevant to the way scientists evaluate informativeness. The information in a theory needs to be extractable in a way that connects with the problems and matters that are of scientific interest.' [Loewer, 2007, pp. 324-5] Similar ideas are expressed in Loewer [fort.].

It's not exactly clear what this notion of 'extractability' consists in, though. In his [1996] Loewer is a bit more specific. He claims we should 'measure the informativeness of a system not in terms of its content (i.e., set of possible worlds excluded) but in terms of the number and variety of its theorems' [p. 186] The discussion that follows makes it clear that 'theorem' is meant in the logical sense — systems are informative, then, in virtue of there being logical derivations starting from the axioms of the system to facts about the world.

The thought is that this conception of informativeness would rule out a system containing only $\forall x F(x)$ from being the best system. From just that axiom and logic there are very few facts about the world that we can derive. From a richer set of axioms, say the laws of Newtonian Mechanics and certain background conditions, we could derive a lot more.

So, switching to this notion of informativeness seems like a sensible option for someone trying to build a package deal account. But, there are difficulties with this approach — perhaps these difficulties are part of the reason that Loewer was less specific about what makes for informativeness in his [2007].

Here is one concern with this approach. Say we have a fact that is informative about a particular predicate P. For us to be able to logically derive that fact from a set of axioms those axioms must already contain the predicate P. This is a problem since the package deal account says that the natural properties are the ones that figure in the axioms of the best system procedure. But, now it seems like just about every property that is involved in the data sentences needs to be part of the axioms of a system if that system is to be informative about the data. Consequently, the best system procedure doesn't do much to pick out the natural properties — the natural properties that result for this view will be given by the properties that the data sentences are about.

How can we deal with this problem? Here's one possible strategy, in the spirit of some comments in Loewer [fort.]. We can suggest that the best system contains both the laws, and some additional principles that do not have the status of laws – *bridge principles* that connect up the laws with other facts. Further, we can say that the natural properties are the properties that are involved in the laws – a property being involved in a bridge principle does not make it natural. This strategy opens up the possibility of there being derivations from the best system to a wide range of data sentences without it being the case that the that the natural properties are just determined by whatever the properties that the data sentences are about.[3]

A strategy along these lines sounds promising. In part because, as we mentioned in passing earlier, standard versions of the BSA already distinguish between the axioms of the best system that count as laws, and the axioms that do not count as laws. Typically this is done to allow the best system to contain background or initial conditions, along with the laws.

Imagine our set of axioms was just the laws of Newtonian mechanics — such axioms aren't really

[3] Siegfried Jaag suggests to me that a defender of the PDA might try a different strategy – claiming that the systemization procedure fixes the data language at the same time as it fixes the laws and natural properties. How exactly this can be done in a way that avoids trivialization is extremely unclear though. If not only the laws and the natural properties but the data itself is determined internally to the PDA then it is hard to see exactly what the problem is with systems like the one that contains only $\forall x F(x)$. After all, such a system is very informative about the facts about F.

very informative unless they are augmented with background conditions about, for example, what objects there are and their initial physical state. So, the standard BSA typically says that the axioms of the best system can sometimes include such background conditions. But clearly, these background conditions don't count as laws.

How might we make this distinction between the axioms of the best system that count as laws and the axioms that do not? As we mentioned in section 1 Lewis's approach to this was syntactic — the laws are the axioms of the best system that have the form of *universal generalizations*. The axioms that had some other logical form count as background conditions. But regardless of how we make this distinction, it seems like we need a distinction between the axioms of the best systems that are laws and the ones that are not.[4]

The strategy that we are considering – one that Loewer [fort.] sometimes seems to accept – is to defend similar distinction in the context of the PDA. That is, to say that in order for a system to be informative it needs to include axioms that are basic laws, but also it needs to include axioms that are bridge principles that connect the vocabulary of the basic laws to the vocabulary of the data sentences.

Such bridge principles, for Loewer, are things like principles that identify 'quartz rock at a particular location in terms of there being a certain arrangement of SiO_4 molecules' and more generally 'principles that underlie connections between fundamental and macroscopic and other non-fundamental sentences'. [Loewer, fort.] The thought then, is that the natural properties are the properties that are part of the axioms that constitute basic laws and not those that constitute bridge principles. Thus we can logically derive the data sentences from the the axioms, while still having the set of natural properties be much narrower than the properties that are involved in the data sentences.

(For Loewer, these bridge principles connect the fundamental vocabulary to facts about macroscopic states because, as we will discuss in the next section, he thinks that the data sentences are about

[4] Interestingly, Loewer and Albert — in developing their influential *Mentaculus* approach — revise the BSA so as to not contain such a distinction. This is in order to allow the *Past Hypothesis* — the claim that the universe started out in a very low entropy initial condition to count as a law.[Loewer, 2012, Albert, 2000]

macroscopic states. If your conception of the data sentences is different that will affect your view of the nature of the bridge principles.)

Lots of questions arise about this strategy though. In particular, how does this version of the PDA distinguish between the axioms that are the basic laws and those that are the bridge principles? Lewis's claim that the laws are the axioms that are universal generalizations doesn't seem to help in this context.[5]

And further, even if we have a way of distinguishing the laws and the bridge principles, there are hard questions about the role of such bridge principles in the systemization procedure. For example, do axioms that constitute the bridge principles count for simplicity to the same degree as the axioms that construe the laws? If they do then there is a concern that there will be far more bridge principles than basic laws. If bridge principles include things like principles that identify a 'quartz rock at a particular location in terms of there being a certain arrangement of SiO4 molecules' then there will have to be similar bridge principles for huge numbers of other objects. And these bridge principles will be very complicated. Consider our current candidates for basic laws — the principles connecting quantum field theory, say, with facts about rocks will be incredibly complicated. So, the simplicity of a system will be almost wholly determined by the simplicity of these bridge principles. But this seems to be the wrong result — the simplicity of the basic laws should matter a lot.

Perhaps, then the bridge principles should not count in measuring the simplicity of a set of axioms. But then it seems that the Predicate F Problem comes back. Imagine a set of axioms where the single basic law is $\forall x F(x)$ but it's combined with lots of other axioms — complicated bridge principles connecting the predicate F to all the other facts that we care about. If the complexity of these bridge principles doesn't matter for judging the simplicity of a set of axioms, then the result is that this set of axioms really is the best system, and the property F is a natural property.

So, for this strategy to succeed, the bridge principles have to count for simplicity, but not to the

[5] Neither, I think does Hall's [2012] way of distinguishing the initial conditions part of a system and the nomic part of a system — his view is that the initial conditions and the laws have a different role with respect to simplicity and informativeness — but getting into the details of that would take us too far astray.

degree that the basic laws do. Exactly how to develop this, and how to distinguish bridge principles from the basic laws is somewhat unclear.

But these problems came from appealing to a conception of informativeness where a set of axioms being informative involves being able to logically derive data sentences from the axioms. Loewer [1996] certainly seems to commit to this. But in his [2007] and [fort.], he is less committal. As we noted earlier he uses a notion of 'extractability' — the axiom $\forall x F(x)$ is not informative because the information contained in that axiom isn't 'extractable'. But, as we noted, it's not clear what this notion of extractability comes to. Loewer doesn't give us a precise account to work with.[6]

I don't think this is an oversight or a failure on Loewer's part, though. Rather, I think, the reason that he doesn't give us much information about the nature of extractability or the way that $\forall x F(x)$ is uninformative is because he wants to rest upon the answers that scientific practice implicitly gives to these questions.

For example, in his [2007] directly after discussing the way in which the axiom $\forall x F(x)$ is uninformative he says that:

> The criteria for evaluating candidate systems are determined with an eye toward their resulting in systems that provide scientifically significant information in forms that are useable to scientists for prediction and explanation. These criteria have been developed and refined during the history of physics since the first proposals for law specifying theories. Especially important is the extent to which a candidate system supports predictions and explanations of fundamental events and regularities, events and regularities of the

[6] In personal correspondence Loewer has suggested that his current view takes informativeness to be something like a priori deducibility, though he's undecided about exactly how this deducibility works. The overall idea is that the PDA fixes on a set of axioms and those axioms are informative when we can a priori deduce facts about the data from those axioms. So, perhaps, we can, given the axioms, a priori deduce facts about properties that are not mentioned in the axioms – thus avoiding the problem under discussion.

This is an extremely interesting suggestion and there is a huge amount to say about which facts are a priori deducible from what base – see, for example, the vast and incredibly complex discussion in Chalmers [2012]. Perhaps the PDA is best developed by construing informativeness as a priori deducibility, but that would require at least a whole other paper to explore. So, for the rest of this paper I'm going to stick to the suggestions that arise in Loewer's written work.

> special sciences and more generally of phenomena that come to be seen as important to the scientific community.

The thought, I take it, is that the precise criteria for judging the bestness of systems, and the answers to questions like the ones we have been raising about bridge principles, have been developed over the years by scientific practice (or, perhaps just the practice of physics). So, there is a sense in which we don't need a philosopher to come up with these criteria, we just need to point to those that are implicit in science. It's clear that scientists would not take $\forall x F(x)$ to be informative — so, that means that there must be some conception of informativeness, implicit in science, that rules out such an axiom. This thought runs through lots of Loewer's work on the PDA, and the work of many other Humeans. We will discuss these ideas in much more detail later.

4 Role (2) and the language of the data

So, Loewer wants to deal with the predicate F problem by switching to a notion of informativeness on which $\forall x F(x)$ does not count as informative. As we noted, there are challenges in making sense of this notion of informativeness. But, for now, let's put this aside and assume that such a notion is available.

Even if we can, by appealing to such a sense of informativeness, avoid the result that trivial systems — like that where the only axiom is $\forall x F(x)$ — count as best there is still a concern about how the PDA could get the right results. Is it really the case that only natural properties will be involved in the axioms of the PDA? (Of course, if the PDA is true then it is a definitional truth that only natural properties will be involved in the regularities of the PDA. The issue is about whether the properties that the PDA says are natural are good candidates for being so.)

In particular, the major concern is that the facts about the world include facts about both intuitively natural and intuitively unnatural properties. It is a fact about the world that this grass is green, but it's also a fact about the world that this grass is grue. In fact, there will be far more facts about intuitively unnatural properties than intuitively natural ones, since the natural properties are a small subset of all the properties. According to the PDA, a system is informative in virtue of entailing facts about the world. But it's hard to see why the intuitively natural properties will be effective at entailing facts about predominately unnatural properties.

The obvious way to deal with this problem is by restricting the data sentences in some way. Someone developing the BSA might restrict the data sentences to only involve natural properties. But clearly this option isn't available to a defender of the PDA. Loewer, in his (1996) and his (2007), suggests two strategies for restricting the data sentences. (Loewer actually never describes exactly why such strategies are required — they are both just integrated into his description of the PDA. But I take it that the reason they are there is, at least in part, to deal with the issue we just described.)

The first strategy involves restricting the phenomena that we are interested in, and thus restricting the data sentences. The second involves restricting the language that the data sentences can be formulated in.

He explicitly appeals to both of these strategies in his (2007) where he says that 'A candidate for a final theory is evaluated with respect to, among the other virtues, the extent to which it is informative and explanatory about truths of scientific interest as formulated in SL or any language SL+ that may succeed SL in the rational development of the science.' [p. 325] ('SL', for Loewer is the language of 'scientific english'.) But let's start by considering the first strategy, restricting the phenomena that the data sentences describe.

4.1 Restricting to Phenomena of Interest

Again, in his (2007) Loewer suggests that the PDA should restrict the data to deal with 'truths of scientific interest'. The notion of scientific interest is a little vague, in his (1996) he has a slightly more specific proposal for how to restrict the data, in the same spirit.

> I assume that it is the job of physics to account for the positions and motions of

> paradigm physical objects (planets, projectiles, particles, etc.). This being so, the proposal is that we measure the informativeness of an axiom system so that a premium is put on its informativeness concerning the positions and motions of paradigm physical objects.[p. 186]

(A similar idea is expressed in Loewer [fort., p. 21].)

Strictly, this doesn't tell us to restrict the data so that the only informativeness we care about is informativeness about paradigm physical objects, rather it tells us to weight informativeness about paradigm physical objects highly, but the basic strategy is the same.

However, restricting the data in this way doesn't seem to do much to solve our problem. Even when we focus on true sentences about paradigm physical objects, there are huge numbers of true sentences about those paradigm physical objects that are stated in non-natural terms. In fact there will be far more true sentences ascribing unnatural properties to paradigm physical objects than those ascribing natural ones. For example, for every sentence expressing that a certain paradigm object has a certain acceleration there will be many others expressing that it has grueified versions of acceleration. Restricting to a certain topic will not exclude unnatural properties from being part of the data.

4.2 Restricting the Language of the Data

More promising is the second strategy — restricting the language that the data sentences are formulated in. Loewer suggests this strategy in his (2007) (in combination with the strategy of the previous section), though not in his (1996). Again, the suggestion is as follows:

> 'Let SL be a present language of science, say scientific english (English supplemented by the languages of mathematics, fundamental physics, and the various special sciences). A candidate for a final theory is evaluated with respect to, among the other virtues, the extent to which it is informative and explanatory about truths of scientific interest as

> formulated in SL or any language SL+ that may succeed SL in the rational development of the science.' [p. 325]

This strategy looks like it will do better at ensuring that the PDA only outputs laws that involve natural properties. The restriction to scientific English rules out facts about gerrymandered, very unnatural properties as being part of the data and so it's far more plausible that the axioms that best balance simplicity and informativeness with respect to this data will be the natural properties.

But the obvious concern with this is that restricting the language of the data seems to undermine the point of the PDA. Notice that the restriction of the data to scientific english privileges a certain set of properties as special — the properties that the predicates of scientific english refer to. But, the aim of the PDA all along was to pick out a set of properties as special. So it's not clear why it is legitimate to start by privileging the properties of scientific english.

In fact, if we start by restricting the data to the language of scientific english then we might start to wonder what the point of the whole systemization procedure is — isn't the restriction to scientific english doing all the work at picking out the right properties?

This, though, is taking the point a bit too far. Notice that there are properties that are referred to by scientific english that are not perfectly natural. For example, it's common in scientific english to talk about measurement, for example, but *being a measurement* isn't a natural property. So even if we restrict the data to be formulated in scientific english the best system procedure does have some work to do.

There is a reasonable version of the view, then, where we start with some privileged language in which the data is formulated — say, the language of scientific English — and then from this starting point the PDA identifies what the natural properties are. That is, the PDA tells us about the properties active in our most basic scientific theories. Of course, this view perhaps doesn't get us as far as we would have hoped — we need to start with some restriction that already rules out lots of gerrymandered, very unnatural properties; we don't generate the natural properties out of nothing. But still, the systemization procedure does do serious work.

The most obvious issue to be worked out with this approach is what exactly this initial language of the data is, and how to justify privileging this language. Loewer says that data should be formulated in scientific English or some development of scientific English, but understood literally this seems very implausible. Why English? Why not Cantonese? Or Telugu? Rather, we should move to some formal language. But then, of course, we still need to restrict the relevant formal language to exclude extremely unnatural properties like the predicate F, or else we have got nowhere. And exactly how we restrict the language, and what the justification for such a restriction is are very hard questions — they mirror the questions that motivate the PDA approach in the first place.[7]

5 Ideal Scientists

It might seem like the defender of the PDA hasn't made much progress then — they once again face the question of how and why to restrict the systemization procedure to some privileged language. But instead of restricting the language of the axioms, like the traditional BSA, the PDA restricts the language of the data.

However, I think this isn't right. The discussion here in fact points to an interesting feature of Loewer's approach, and of the approach of many modern Humeans.

As we started discussing at the end of section 3.1 there is a certain way in which the details we have been considering in the last two sections — of what exactly the language of the data is, of how exactly to define informativeness, of what the role of bridge principles are, and so on — don't really get to the heart of the matter, at least as Loewer develops it in his (2007) and (fort.). Loewer's guiding thought, I think, is a kind of ideal observer view of naturalness — the natural properties are whatever an ideal scientist would use in their most basic physical theories.

[7] In personal correspondence Loewer suggests another option – that instead of fixing on one particular privileged language of the data rather the data can be in any language, but the criteria which govern the bestness of systems should be invariant across translations. The thought is that actual science takes place in multiple languages, but this doesn't affect scientist's judgements about which system is best. Unless we give a substantial account of those criteria and how they can be invariant across translation then this strategy involves appealing to the criteria that are implicit in science to solve the problem. I'll discuss this kind of strategy in the next section.

An ideal scientist, here, isn't a scientist who just believes all and only the true theories. Rather, it's someone who ideally implements the methods that are implicit in actual science. So, take the methods for deciding what the fundamental laws are that are used in actual scientific practice, and imagine that those methods and their implementation are idealized in certain ways — we idealize away from our ignorance of the non-nomic facts, our limitations on computational power, and so on. On this approach the natural properties just are the properties that a scientist ideally implementing these methods would use in their most basic physical theories.

One way to give this view some intuitive support is to note that science seems to be a very good guide to what the natural properties are. Good candidates for natural properties are things like *mass*, *charge* and *spin* – properties that play central roles in our basic science. So an initial idea is just to define the natural properties as whatever scientists say that they are. But this is too simple, because there are many ways that scientists could be in error. They could be missing important observations that would be relevant to their theories, they could make logical or mathematical errors, they could just fail to think of relevant possibilities, and so on. So, imagine a counterfactual scientist who didn't make such errors and had all the relevant information — such a scientist would, presumably, be recognized by current scientists as being in an improved epistemic situation. The thought is that the natural properties are whatever this counterfactual scientist says they are.

This kind of ideal observer view is described clearly in Hall [2012], and seems to be the guiding idea for other modern Humeans (e.g Jaag and Loew [fort.], Dorst [2019]).[8]

In fact, the traditional BSA is very commonly described in a similar way – as taking the epistemology of laws and raising the principles implicit in this epistemology to an account of the metaphysics of laws. Simplicity and informativeness are clearly relevant considerations for scientists picking out the laws — the traditional BSA makes them constitutive of the laws. For example, Carroll [1994, p.45] says that defenders of the BSA '[let] their metaphysics be shaped by the epistemology of law-

[8]Though there is interesting debate about exactly how idealized this ideal observer is. Hall, for example, calls his ideal observer a 'limited oracular perfect physicist' and a theme in some of the work just cited is that we should understand the ideal observer to be limited and to be more 'like us' in important ways. [cite stuff from this volume?]

hood'. And Hall [2012] says that 'our implicit scientific standards for judging lawhood are in fact constitutive of lawhood'.

And further, as Lewis himself notes, the epistemology of laws and naturalness are very closely connected. In fact, it is from Lewis's description of the epistemology of laws and naturalness that Loewer gets the phrase 'package deal':

> Thus my account explains…why the scientific investigation of laws and of natural properties is a package deal; why physicists posit natural properties such as the quark colours in order to posit the laws in which those properties figure, so that laws and natural properties get discovered together. [Lewis, 1983a, p.368]

Although Lewis's account of the metaphysics of laws starts with a set of privileged natural properties his account of the epistemology of laws is very different. The natural properties get discovered along with the laws, as part of the ordinary scientific methods for building theories.

So, if the epistemology of laws and naturalness is a package deal, and the BSA involves taking the epistemology of laws and raising it to the level of metaphysics, then it can seem obvious that the Humean could, and should, do the same with naturalness, giving a package deal account of the metaphysics of laws and naturalness.

That this ideal observer approach is Loewer's approach is, I think, pretty clear in his (2007) and (fort.), for example, from the weight he puts on the 'rational development of science' in his (2007), and the way that in (fort.) he stresses that the the criteria for choosing what counts as the best system are 'can be gleaned from an examination of practice in physics and the special sciences. As sciences develop these criteria may evolve and new ones develop'.

Given that this is Loewer's strategy, then a detailed and comprehensive account of the criteria for picking out the best system might not seem required. Sure, it would be nice to have an account of exactly the methods via which scientists identify the basic theory, and hence the natural properties, but even if we don't we can still make a substantial metaphysical claim about the natural properties

— that the natural properties are those that would result from an ideal implementation of scientific methods, whatever they happen to be.

In fact, in his (fort.) he lists nine (!) criteria that are relevant to picking out the best system, including things like 'Many sub-systems can be treated as almost isolated so that in typical circumstances the laws apply to them neglecting their environments' and '[the system] enables predictions, explanations and understanding of a wide variety of phenomena via systematic perspicuous principles that connect fundamental to non-fundamental descriptions.' Further, he notes that this isn't 'anything like a complete list of the criteria operative in fundamental physics for evaluating candidate systems or an account of how to balance them'. And 'All [of the criteria] are in need of clarification and elaboration'. It's pretty clear that Loewer doesn't intend to commit to anything particular specific about the criteria for judging systems.

This is very different from Lewis's version of the BSA which clearly committed to strength and simplicity — and to particular precisifications of these notions — as being the criteria by which we judge bestness of systems.

There are two pretty substantial problems with this kind of ideal observer strategy though. Firstly, the strategy leaves it unclear why we should care about the laws and the natural properties. Sure, there is this particular practice of scientific investigation into the basic laws and natural properties. And there is an ideal implementation of the methods implicit in this practice. But why should we care about what those methods output? After all, there is a practice of astrology, and, presumably, an ideal implementation of those methods, but I take it we have no particular reason to care about that. So what is the difference in the case of science?

Perhaps the most obvious suggestion is that the scientific methods are actually effective at getting at the facts about the laws and natural properties, while astrology is not. But, notice that this move is not available to the the defender of the ideal observer view, because for them there is no realm of laws and natural properties that are independent of the methods of science. So it's true that the scientific methods are effective at getting at the laws and natural properties while astrology is not, but that's only by stipulation. A defender of astrology could equally define laws and natural properties as the output of the ideal implementation of those methods, and so claim that astrology is is effective at getting at the laws and natural properties while science is not.

So, there needs to be some other justification for why we care about the output of the ideal implementation of scientific methods for discovering the laws and natural properties. And, of course, there are options available. Perhaps we care about the ideal implementation of scientific methods because the laws and natural properties outputted by those methods are pragmatically useful to us — they help us achieve our epistemic and non-epistemic goals.

But, in order to give this alternative story for why we should care about the ideal implementation of scientific methods, we need to know what those scientific methods actually are. In order to argue that those methods output things that are pragmatically useful to us we must look in detail at the actual methods — we cannot, remain uncommittal about those methods in the way Loewer does.

Why is this? Why can't we remain uncommittal? Can't we just argue that we know that actual science is extremely pragmatically useful, even if we can't specify exactly what methods it uses. So, we have reason to think that the ideal implementation of those methods will output things that are pragmatically useful.

When considering this line of thought it's important to remember that what is currently under investigation are the basic scientific laws and the natural properties. That is what the PDA is supposed to output — the laws and properties of our most fundamental science. So, the question under consideration is why we should care about the ideal implementation of scientific methods for discovering these basic laws and natural properties.

So, restricting to our basic science, the line of thought is that out actual basic science is extremely pragmatically useful so we have reason to think that the ideal implementation of the methods for discovering these basic laws and natural properties are pragmatically useful.

However, it's just extremely unclear whether the actual laws and natural properties of our basic science are particularly pragmatically useful to us. We just cannot use our basic scientific theories to make predictions about, or manipulate, most of the things that we care about. It's not quantum field theory that we use to work out how to make an airplane fly, or to perform heart surgery.

Of course, science as a whole is extremely pragmatically useful to us. It outputs lots of extremely useful heuristics. And higher-level sciences, like thermodynamics, chemistry, biology and so on, have great pragmatic value. But it's very unclear whether our basic science is especially pragmatically useful. Consequently, it's rather doubtful that the ideal implementation of our methods for outputting the basic laws and natural properties will be pragmatically useful. (Perhaps someone might argue that our current basic science is not especially pragmatically useful, but our ideal basic science will be. Maybe this is true, maybe not – I don't see much evidence one way or the other about what would happen in this situation that is very far from the actual world.)

Consequently, I don't think an appeal to the utility of current science can help us much in establishing why we should care about the ideal implementation of the scientific methods for identifying the laws and natural properties of our basic sciences.

In order say why we should care about the ideal implementation of these scientific methods we need more details about exactly what they are and how they work. If we have such details then, perhaps, we can give a pragmatic story about why we should care about then.

But further, once we have done this — once we have a pragmatic story for why we care about the output of ideal scientific methods — then there is an additional problem: the ideal observer view seems to lose it's force. Imagine a situation where the methods that output things which are pragmatically useful to us and the ones that an ideal scientist would implement are different. If the reason that we care about the methods of science is their pragmatic utility, then in such a situation we shouldn't care about what the ideal scientist says. Once we have given this pragmatic story then, this is what seems to do all the work, and the appeal to the ideal observer drops out. Similar reasoning applies to any alternative story we might give for why we should care about the output of the ideal scientist.

The moral here is that we need to be committal — we need to give a substantial account of what the criteria for bestness of a system are and why we should care about what those criteria output. If we instead just rely on whatever an ideal scientist would say then it's hard to see why we should care.[9]

The second concern with the ideal observer strategy is perhaps more direct. The idea that we can give a reductive account of the natural properties as part of the output of the ideal implementation of scientific methods assumes that we can characterize those methods without already assuming facts about naturalness. I think this is rather doubtful though.

For example, one of the criteria Loewer (fort.) lists as being important in the scientific practice of deciding on a best theory is that the theory 'enables predictions, explanations and understanding of a wide variety of phenomena via systematic perspicuous principles that connect fundamental to non-fundamental descriptions.'. This seems right – our fundamental theory should have explanatory power over a wide range of phenomena, both fundamental and non-fundamental. But, the concern is that in order to recognize whether a theory provides explanations and understanding we already need some grip on the distinction between natural and unnatural properties.

In particular, many putative explanations are bad explanations, and do not yield understanding, precisely because the properties involved are unnatural. For example, take some particular event E that occurs this year. Imagine that the laws are deterministic so that there is some set of possible initial conditions of the universe $\{i_1, \ldots, i_n\}$, such that E occurs if and only if one of these initial conditions held. Then consider some property P that holds of the universe if and only if the initial conditions of the universe are one of $i_1, \ldots, i_n$. Given this, the fact that the universe has property P seems like it might explain the event E. After all, the event E holds if and only if the universe has property P. And the universe having property P leads to the event E occurring. Nevertheless, it's clear

[9] Notice that the way an ideal scientist was characterized in this discussion was as someone who ideally implements certain scientific methods. This fits with the general strategy discussed in this section of taking the methodology of actual science and raising it to a metaphysical account of laws and natural properties.

We could characterize what an ideal scientist is very differently – for example, perhaps an ideal scientist is someone who generates theories that are maximally pragmatically useful to us. Clearly if this is what an ideal scientist is, there is no problem in explaining why we should care about what they output. But, if we characterize ideal scientists in this way it's not at all clear what the connection is between such an imagined scientist and the actual practice of science. Loewer's consistent appeals to the actual practice of science as being relevant to how the PDA works would be somewhat puzzling on this approach.

that this is a very bad explanation and that's because the property P is unnatural – it is just not the type of property that is appropriate for explaining.

Cases like this suggest that we can't characterize explanatory power independent of judgments of naturalness. And so if the ideal scientific methods for deciding on a best theory involve judgments of explanatory power then they cannot be characterized independent of judgements of naturalness.

More generally, as we noted at the start of the paper, the literature on naturalness has argued that the concept of naturalness is involved in a huge range of other concepts like similarity, intrinsicality, induction, causation, explanation, reference and so on. When we see the wide-ranging roles that naturalness seems to play then it becomes rather implausible that we could characterize ideal scientific methods without mentioning things that presuppose naturalness.

Of course, it is open for the defender of the this approach to the PDA to deny that naturalness does, in fact, play this role with respect to explanation. But then they need to give an alternative story about the badness of explanations involving the property P. And similarly, they could deny that deny that naturalness plays a role with respect to similarity, induction, causation etc. giving a story about those concepts that doesn't presuppose naturalness. This seems like a very hard task. It's seriously hard to see how to, for example, give a characterization of scientific practices of explanation and induction without already presupposing exactly the thing you are trying to give an account of in the PDA – that is without presupposing a distinction between the properties that it's legitimate for science to theorize in terms of and those that are not legitimate for science to theorize in terms of.

For these two reasons, the ideal observer strategy is, I think, not successful. More plausible is a strategy which is more committal — which really specifies the criteria that make theories best. Though as we noted in sections 4.1 and 4.2 there are very big challenges in properly developing this strategy. But, I suspect, this is the path the defender of the PDA should take.

6 Role (i) and the Metaphysical Underpinnings of the PDA

Even if such challenges are met there is another way in which the PDA faces challenges that do not face the traditional BSA. In particular, Lewis's version of the BSA was built upon the Humean mosaic — the mosaic, as we discussed earlier, is the distribution of perfectly natural properties that are instantiated at spacetime points. The mosaic provides the fundamental basis on which the BSA is built. But what is the the metaphysical picture underlying the PDA?

On the PDA laws and natural properties are clearly metaphysically non-fundamental, but what things are metaphysically fundamental? Perhaps the best answer for the defender of the PDA is agnosticism — the details of the underlying metaphysics doesn't matter so long as there is enough structure to act as truth makers for the data sentences. This answer is suggested by some parts of Loewer (fort).

This approach does mean, however, that a lot of the traditional motivations for the BSA don't apply to the PDA. The BSA was, for Lewis, part of a project of defending a Humean view about the metaphysical basis of the world — that there are no necessary connections between distinct existences; the world is just one thing after another. The BSA gave an account of laws that followed from that view of the metaphysics.

But a PDA that is agnostic about the underlying metaphysics is not clearly a Humean view. The fundamental metaphysics might contain necessary connections, or it might not. Perhaps, in order to retain its Humean character, the PDA should not be completely agnostic about the underlying metaphysics — perhaps defenders should say that the world, at the fundamental level, does lack necessary connections. But this raises another concern: If the underlying metaphysics of the world is Humean – if, for example, it consists in just certain properties being instantiated at spacetime points with no necessary connections then doesn't this metaphysics just provide us with a set of privileged properties? Shouldn't we just say, as Lewis does, that the natural properties are the ones that make up this underlying metaphysics?

The defender of the PDA view has to say that the answer to this question is no. Even if underlying metaphysics privileges a set of properties those are not the *natural* properties — they are not the properties that are part of accounts of intrinsicality, laws, induction, causation, explanation, reference and so on. Rather, it is the properties that are the output of the PDA that have these connections. This view is, of course, defensible — in fact, it's in the spirit of Dasgupta's [2018] argument (that Loewer (fort.) mentions) that even if there were some primitive natural properties they should not guide our scientific theorizing — but it's another substantial commitment for the defender of the PDA, that the best candidate for the natural properties is the output of the PDA and not the the properties that make up the underlying metaphysics.

7 Conclusion

In the traditional Humean picture the notion of naturalness plays multiple roles — as part of the mechanics of the BSA and as giving an account of the underlying metaphysics. The PDA, in aiming to integrate a reductive account of naturalness with a reductive account of laws, has to develop an approach where naturalness doesn't play these roles. Consequently, the PDA is not just a small adaption of the traditional Humean BSA, it's a radically different, and novel, approach.

This approach faces many difficulties that the traditional BSA does not. However I do think there are versions of the PDA that might be defensible. In particular, a version that is (i) agnostic about the underlying metaphysics and (ii) heavily restricts the language of the data sentences, might be feasible. As we have noted, though, such an account still faces significant challenges to overcome. But perhaps more importantly, even if successful, this account is just far more restricted than we might have originally hoped. Such an account starts by privileging some particular language, and thus some particular properties — part of the work that we might have hope the PDA would do has to be done before the account gets going.

In light of these problems you might naturally want to rest more heavily on the methods implicit in science as outputting the natural properties. But, I argued, this ideal observer view is not particularly promising.

One last suggestion: Maybe the PDA is a good tool for a slightly different task than the one we have been considering so far. Loewer uses the PDA to give an account of the natural properties, where the natural properties are construed as having a very close relation to the properties of fundamental physics. This is in line with most of the literature on naturalness. But, we can make a similar distinction between the natural and unnatural properties in the special sciences too. And it's not easy to see how this special science naturalness could arise from an account of fundamental-level naturalness. Perhaps the PDA could be used as an account of this higher-level, special science naturalness and not as an account of fundamental-level naturalness. So, for example, the laws and natural properties of biology could, perhaps, arise from an application of the PDA where the data is formulated in terms of the perfectly natural properties and restricted to the paradigm phenomena of biology. I suspect that building a PDA of this kind, upon a base where we assume some fundamental notion of naturalness, allows the view to be developed more smoothly and compellingly. But that's a discussion for another time.

References

David Z Albert. *Time and Chance*. Harvard University Press, 2000.

Karen Bennett. *Making Things Up*. Oxford University Press, 2017.

John Carroll. *Laws of Nature*. Cambridge University Press, 1994.

David Chalmers. *Constructing the World*. Oxford University Press, 2012.

J Cohen and C Callender. A better best system account of lawhood. *Philosophical Studies*, 2009.

Shamik Dasgupta. Realism and the absence of value. *Philosophical Review*, 127(3):279–322, 2018.

Heather Demarest. Powerful properties, powerless laws. In Jonathan D. Jacobs, editor, *Causal Powers*, pages 38–53. Oxford, United Kingdom: Oxford University Press, 2017.

C Dorr and J Hawthorne. Naturalness. *Oxford studies in metaphysics*, 8, 2013.

Chris Dorst. Towards a Best Predictive System Account of Laws of Nature. *The British Journal for the Philosophy of Science*, 70(3):877–900, 2019.

J Earman and J Roberts. Contact with the nomic: A challenge for Deniers of Humean Supervenience about laws of nature part II: The epistemological argument for Humean Supervenience. *Philosophy and Phenomenological ...*, 2005.

M. Eddon and Christopher J. G. Meacham. No work for a theory of universals. In Jonathan Schaffer and Barry Loewer, editors, *A Companion to David Lewis*, pages 116–137. Wiley-Blackwell, 2015.

N Hall. Humean Reductionism About Laws Of Nature. *Unpublished manuscript*, 2012. URL `http://philpapers.org/archive/HALHRA.1.pdf`.

Siegfried Jaag and Christian Loew. Making best systems best for us. *Synthese*, pages 1–26, fort.

David Lewis. New work for a theory of universals. *Australasian Journal of Philosophy*, 61(4):343–377, 1983a. doi: 10.1080/00048408312341131.

David Lewis. *Philosophical Papers, Volume 1*. Philosophical Papers, Volume 1. Oxford University Press, Oxford, 1983b.

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

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

Barry Loewer. The emergence of time's arrows and special science laws from physics. *Interface Focus*, 2(1):13–19, Feb 2012.

Barry Loewer. The package deal account of laws and properties. *Synthese*, fort.

Tim Maudlin. *The Metaphysics Within Physics*. The Metaphysics Within Physics. Oxford University Press, New York, 2007.

Theodore Sider. *Writing the Book of the World*. Oxford University Press, 2011.

Tuomas E. Tahko. Fundamentality. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, fall 2018 edition, 2018.

Bas van Fraassen. *Laws and Symmetry*. OUP, 1989.