Towards a Theory of Special Science Naturalness

Harjit Bhogal

In recent decades, particularly since the publication of Lewis's 'New Work for a Theory of Universals' (1983), it has become common for philosophers to recognize the need for a distinction between *natural* and *unnatural* properties. Natural properties, it is thought, play a variety of important theoretical roles — for example, to do with laws, induction, and similarity, amongst others.

This literature on natural properties has developed in a way that typically closely associates the natural properties with the properties of fundamental physics. Candidate natural properties are things like *spin* and *charge*. Further, the natural properties are generally taken to be metaphysically fundamental, unlike the properties of special sciences, which are metaphysically derivative (see, for example, Dorr and Hawthorne (2013)).

However, there is a similarly compelling theoretical need for a notion of naturalness that applies to the properties of the special sciences as there is for one that applies to fundamental physics. Just as there are certain, natural, properties that play central roles in the practice of physics with respect to physical laws, physical explanations and fundamental level similarity (amongst other things) so there are certain natural special science properties which play analogous roles with respect to the practice of the special sciences. In this paper I start to develop an account of such special science naturalness.

(An important clarification: My project is about 'naturalness' in the sense of the literature following Lewis (1983). There is another common sense of naturalness — the naturalness of *natural kinds*. This literature focuses more on issues of *classification* — how to classify organisms into species, for example. These two issues are not totally distinct, but it's important to be clear that my project is about natural properties, not natural kinds.)

Notably, the account that I'm going to discuss is unlike traditional accounts of fundamental naturalness. (From now on, I'm going to use the term 'F-natural' (for 'fundamental natural') for the traditional, non-special-science natural properties. Where there is no ambiguity I will use 'natural' to mean 'special science natural' but never to mean F-natural.) On this account special science naturalness is not a primitive — it's not some irreducible intrinsic property that other properties have. Rather, the account is reductive — and further, the account integrates naturalness with the practice of the special science. So, such a story about special science naturalness can be attractive to people who wholeheartedly reject primitivist accounts of F-naturalness, or even the notion of F-naturalness itself.

In the first section I look at some of the roles special science natural properties play. In section 2 I briefly consider some previous accounts of special science naturalness. In section 3 I motivate the core of my approach before developing it in more detail in section 4. In section 5 I look at how the approach of section 4 fits with our intuitive judgments about what properties are natural and I discuss a limitation to the account.

Because of this limitation the account doesn't yield a full set of necessary and sufficient conditions on special science naturalness, but, I think, it goes a long way towards elucidating its nature.

1 New Work for a Theory of Special Science Naturalness

I'm going to start, in this section, by briefly considering the roles that special science natural properties play – for example, in connection with laws, causation, and explanation. Just as in the literature on F-naturalness, these roles both motivate the need for a notion of special science naturalness and help us identify which properties are natural.

Laws: Special science laws are typically understood to involve natural properties. Fodor (1974, p. 102) makes this point when he says: 'I take it that there is no natural law which applies to events in virtue of their being instantiations of the property *is transported to a distance of less than three miles from the Eiffel Tower*'. In fact, he goes on to suggest that the natural properties are exactly those that are part of some law. (See also Lewis (1983), Armstrong (1983) and Loewer (1996), amongst many others, for other views of laws that accept a tight connection between naturalness and laws. Cohen and Callender (2009) are a rare exception in that they develop a view of laws that explicitly allows unnatural properties to be part of laws.)

Similarity: Worlds can be similar in many respects while being radically different at the fundamental level. For example, two worlds can be similar in virtue of both containing market economies, even if the underlying physics of the two worlds is extremely different. One world, for example, might be made up of particles while in the other the basic building blocks are fields, or strings, or something even more exotic.

Given this difference in the physics, the similarity between the worlds does not seem to be grounded in any sharing of F-natural properties. Rather, the two worlds are similar because they share certain high-level properties.

But clearly, only some high-level properties can ground similarity. If every high-level property could ground similarity then I would be similar to the Eiffel tower in virtue of us sharing the property *being a philosopher or taller than 900 feet*. This is a job for special science naturalness; the natural properties are the ones that ground similarity, the unnatural ones do not.

Causation: Lewis (1983, p. 368-9) points out that if naturalness is required for laws and for similarity, then it looks like it is required for causation too, given that the most popular theories of causation are either nomic or based on counterfactuals that appeal to the notion of similarity.

Further, he points out (p. 369-370) that a distinction between natural and unnatural properties is required to avoid implausibly multiplying the number of causes for every event. If an event is caused by the instantiation of a particular property — that property instantiation is relevantly nomically or counterfactually connected to the event — then there will be many other instantiations of intuitively unnatural properties that that are similarly connected to the event.

1.1 Explanations

Most importantly for what is to come, however, is the connection of naturalness to explanation. In particular, special science naturalness is needed to rule out explanations that seem inappropriately gerrymandered. (Others who discuss this problem include Kitcher (1981, section 8), Strevens (2008, section 3.6.1), Franklin-Hall (2016) and Weslake (2015).)

Consider the explanation which cites the fact that the ice cube was dropped in warm water to explain it's melting. Or the explanation that cites the introduction of a predator to explain the extinction of a species. Or the explanation that cites an increase in demand to explain an increase in price. All of these seem like they have the right form for an explanation. They would be good explanations if the world cooperates; for example, if demand really did rise and that fed through to rising prices.

But some putative explanations do not have this appropriate form. Consider the fact e about the price of a particular good rising. Assuming that the fundamental laws are deterministic[1] then there are certain physical states of the world $b_1, \ldots b_n$ such that e holds if and only if one of those states held at time t. And then consider a property P that the world possesses if and only if it is in one of the states $b_1, \ldots b_n$. Now consider the putative explantation of e from the fact that the world has property P at time t. This is not a good explanation of e. In fact, it doesn't really look like an explanation at all.

What is wrong with such an explanation? It seems like there is something wrong with the property P that makes it inappropriate to appeal to in explaining other facts. Again, this looks like a job for special science naturalness – P can't be used to explain because it is unnatural.[2]

In section 3 I'm going to look more closely at this problem in order to motivate my approach to explanatory correctness.

[1] We will assume determinism throughout what is to come.

[2] One might be tempted to think that that we don't need a complicated account of naturalness here — the reason that P cannot be used to explain is because it is a *disjunctive property*. It's not clear, though, in what sense P is a disjunctive property that does not also apply to paradigm special science properties that we do want to use in explanations — properties like *being a predator*. For example, both properties are disjunctively defined in terms of the F-natural properties and Clapp's (2001) influential definition of disjunctive properties does not distinguish these two properties.

2 Theories of Special Science Naturalness

Before we move on to that, though, we should quickly consider some alternative accounts of special science naturalness. I certainly won't have enough time to fully survey the literature, or even to give a detailed discussion of the accounts that I do consider. My aim is simply to help to locate my approach within the existing literature.

2.1 Primitive Special Science Naturalness

One option is to take special science naturalness as primitive, just as most of the literature on F-naturalness takes that as primitive. Loewer (2009, especially p. 228) claims that Fodor (1974) held this view. Some parts of Fodor's paper certainly suggest that, for example:

> I am suggesting, roughly, that there are special sciences not because of the nature of our epistemic relation to the world, but because of the way the world is put together: not all natural kinds…are, or correspond to, physical natural kinds. (p. 113)

One plausible way of interpreting this, is as saying that special science natural properties are just part of 'the way the world is put together', and their naturalness doesn't reduce to anything else. (Loewer understands Fodor in this way.)[3]

Whether or not Fodor actually held this view, I take such a view to be highly implausible. In particular, the view implies that there are scientific facts that are not determined by the facts about the fundamental physics. For example, the facts about what the correct biological explanations are depend upon which biological properties are natural, so on this account, they

[3] Fodor uses the terminology of natural *kinds* but the rest of the paper suggests that he is talking about natural properties in close to the sense that I have been discussing.

would not be determined by the physical facts. Thus the view is in conflict with a plausible version of physicalism. This conflict is enough to motivate a search for an alternative view.

2.2 Pragmatic accounts

On the other end of the spectrum from primitivist accounts are pragmatic accounts. The basic idea is that natural properties are the ones that are most useful for us – the properties that best help us achieve our ends or satisfy our interests (see, for example, Kitcher, 2001).

My primary objection to such accounts is the obvious one – it makes the facts about naturalness inappropriately depend upon our interests. Perhaps such a dependence is acceptable if we have no other options, but it is something that we should try to avoid.

There are a variety of other objections to pragmatic accounts. I'm just going to point to one here. If we want to use our account of natural properties in an explanation of certain features of scientific practice then there is a risk that grounding the facts about naturalness in facts about what we, or scientists, are interested in will make the explanation uninformative, or at worst circular. In this respect pragmatic accounts are less explanatorily powerful than other accounts.

I want to see to what extent I can ground naturalness in objective, non-pragmatic, features of the world. Perhaps pragmatic considerations are something we would have to fall back on if such an account fails.

2.3 Graded F-Naturalness

One existing account that avoids both primitivism and pragmatism involves defining a measure of graded naturalness in terms of the F-natural properties and then claiming that intuitively natural special science properties turn out to be fairly natural on this measure. Lewis (1986b, p. 62) took this approach. He claimed that the degree of naturalness of a property is determined by the length of its definition in terms of the perfectly F-natural properties. Sider (2011, section 7.3) also accepts (a variant of) this approach. In fact this, I take it, is the default approach to the problem of assigning naturalness to special science properties when it arises in the literature on F-naturalness.

There variety of problems, however, with using this approach to identify the special science natural properties. Sider (1995, pp. 363-4) and Hawthorne (2007, pp. 433-4) both argue that many special science properties, both intuitively natural and intuitively unnatural ones, have may infinitely long definitions in terms of the natural properties, and so would all count as equally natural to each other. But there are other problems too.

Firstly, even if special science properties have finite definitions, it is not at all clear that the intuitively natural properties that have shorter definitions than the intuitively unnatural ones. Consider, for example, the economic property of *inflation*. The definition of this property would be incredibly long and we have very little idea about how it would go. The thought that our actual concept of inflation will turn out to have a relatively simple definition compared to intuitively unnatural economic properties, and similarly for other intuitively natural special science properties, just seems hopelessly optimistic.

Secondly, disjunctions (even long disjunctions) of the properties of, for example, chemistry, will have shorter definitions than the natural properties of higher-level sciences, like economics. For example, the property of *being Lithium or Argon or Copernicium or Fluoride or Gold or Polonium or Sulphur* will have a much shorter definition in terms of the F-natural properties than *inflation*. But this will not do for our ends. It is not ok to use intuitively unnatural chemical properties like this in explanations, but it is ok to use economic properties.

Thirdly, it is highly implausible that when we judge certain special science properties to be natural and others to be unnatural what we are responsive to is the length of the definition of the property in terms of the properties of fundamental physics. Again, we have almost no idea how such a definition would go. As such, this account doesn't make sense of our judgments regarding naturalness.

Perhaps these problems can be overcome, say with strategic appeal to functional properties. I'm not optimistic for the success of this project, but I can't discuss that more here. But the difficulties for the graded F-naturalness approach at least motivate us to look elsewhere.

Importantly, all this is not to say that there is no useful concept of graded naturalness. I'll discuss this more later. The point here is simply that it's hard to see how notion of graded F-naturalness can give us an account of the special science natural properties.

3 Motivating the View

In this next few sections I'm going to develop my approach to special science natural properties. In this section I'll motivate the intuitive idea, by looking carefully at the connection between naturalness and explanation. In the next section I'll give a more precise formulation. In section 5 I'll consider how effective the approach is.

In section 1.1 we discussed the problem of disjunctive explanations. In particular, we considered putative explanations of a high-level fact – the economic fact e about the price of good G rising. The problem was that there are certain properties that seem inappropriate to use in explanations of e – for example, the property P that holds of the world just in case it is in one of the states $b_1, \ldots, b_n$, where $b_1, \ldots, b_n$ are all and only the states of the world at an earlier time t that would deterministically lead to e.

P can't be used in explanations because it is an unnatural property. So, if we can find out what is wrong with putative explanations involving P, independent of any thoughts about naturalness, then this should give us a better understanding of the unnaturalness of P, and other properties.[4]

So, what is wrong with the explanation of e from the fact that the world has property P at t? Here is a very intuitive thought: the explanation is bad because P is *gerrymandered*. Talk of 'gerrymandered properties' is common in discussions of naturalness, but it is not always clear what it means. When I say P is gerrymandered, I'm not saying that P is overly complicated — it's not clear in what sense P is more complicated than any other high-level property. And I'm not using 'gerrymandered' as a synonym for 'unnatural' — clearly doing that would not advance the discussion.

Rather, to speak loosely for a moment, what makes P gerrymandered is that it is 'designed'

[4]Again, the issue I'm considering is to do with when and why it is inappropriate to use certain properties, like P, in explanations. There is a related issue that I'm not going to focus on in what's to come. To see this, consider again the putative explanation of e from the fact that (i) the world is P at t, and then consider the putative explanation from the fact (ii) that the world is in b_1, or ...or b_n at time t. Further, assume each of the states b_1 to b_n can be defined in F-natural terms.

There are two reactions you could have to the explanation from (ii). The first is that it is identical to the explanation that cites (i), and so it involves the property P is just the same way (i) does. If this is the case then there is no extra issue, the explanation from (ii) poses the same problems as the explanation from (i) and can be dealt with in the same way.

The second reaction is that, unlike the explanation from (i), the explanation from (ii) does not involve unnatural properties. Rather, the explanation from (ii) involves only natural properties but these properties are combined into a very disjunctive fact. If this is the case then clearly my project of distinguishing between natural and unnatural properties will not do anything to rule out such a putative explanation. So, there is a distinct issue about these kinds of putative explanations, independent of my focus on the connection between natural properties and explanation. (In the few cases where people discuss something like a naturalness constraint on explanations this issue has generally been ignored [Franklin-Hall (2016, p.21), Weslake (2015)].)

Ruling out these kinds of explanations is not especially problematic though. There is a feature of the explanation from (ii) that we can clearly identify as being inappropriate (even if we don't have a sharp theory of why it is inappropriate) — the fact (ii) is hugely disjunctive. (Notice that this is very different from the issue that I am most focused on, because we don't have any such clarity about how explanations involving P differ from good special science explanations.) So, as well saying that genuine explanations cannot involve extremely unnatural properties we should add that explanations cannot involve extremely disjunctive facts.

solely in order to get one, very specific, result — it is designed to be part of an explanation of the fact e — and further it seems to only have this one use. Subsequently, the result that P is designed to achieve seems, in some sense, cheap.

Consider the electoral district rezoning that spawned the word 'gerrymandering'. In that case, the borders of electoral districts were drawn in order to achieve very specific electoral ends. For example, a party might pack most of the supporters of their opposition into a few districts so they lose those districts by a very large amounts but win the majority of districts by much smaller amounts. In paradigm cases of gerrymandering the electoral districts are drawn such that *the only reason* anyone would draw the boundaries in that way is to achieve this very specific electoral goal.

One possible aim of drawing district boundaries a certain way is to have some homogeneity of preferences among people in each district, so that one representative can reasonably represent them all. Another aim is to have districts be geographically compact so that residents of each district can easily interact with each other. There are many more potential aims of rezoning (see, for example, Butler and Cain (1992)). In paradigm cases of gerrymandering we ignore these aims of rezoning, instead focusing on achieving the electoral goal. There is no reason to draw the boundaries in the way we do independent of the electoral aim. Subsequently, the electoral result seems cheap and unearned.

P looks to be gerrymandered in a similar way — there is no reason to appeal to P in our scientific practice independent of way that it is specifically targeted at e. The property P is 'designed' solely to play a role in the putative explanation of e from the fact that the world is P at time t. This gerrymandering makes the use of P in explanations seem inappropriate.

I think this intuition is forceful, but it is rather blunt. In particular, talk of properties being 'designed' for certain ends is rather imprecise. Don't we always design our explanations to fit the phenomenon under investigation? And in what sense do we design properties?

So we need to be a bit more careful here. I'm going to discuss two ways in which P seems to be gerrymandered where typical special science properties are not.

3.1 Portability

The first point is that P is not *portable* — it is not a property that is useful in explaining a wide range of explananda. On the contrary, P only appears to be useful in very few explanatory situations. If we are interested in explaining e, then, as we noted, we might want to appeal to P in the explanation.[5] And perhaps we might want to appeal to P when explaining facts that are extremely closely connected to e. But if we want to explain some other economic fact — for example, a fact about the marginal cost of producing good G — then appealing to P looks like a bad strategy. And further, there don't seem to be biological or chemical or sociological contexts in which we would appeal to P in order to explain. In fact, in the vast majority of cases P does not seem useful to be at all in explaining what we want to explain.

Genuine scientific explanatory practice typically involves explaining a wide range of facts by appealing to the same set of properties. For example, it was considered a deep benefit of Newtonian mechanics that it could explain such a wide range of phenomena, both celestial and terrestrial, starting from a sparse base of properties — force, mass, velocity, and so on (Kitcher, 1981, section 3). Similarly, it seems like a good feature of standard microeconomics that it can explain so much in terms of a few basic properties — properties like

[5] Of course, ultimately we would not want to use P in the explanation of e, because P is unnatural and gerrymandered. But since our task now is to get a grip on what makes P unnatural and gerrymandered it's dialectically appropriate to ignore this for now.

agent, *goods*, *preference*, *demand*, *supply*, *factor of production*, *choice*, *price*.[6] We do not find non-portable properties like P being used in genuine explanatory practice.

So, the use of non-portable properties in explanation is at odds with scientific practice. And there is a good reason for this, because appealing to non-portable properties to explain seems to leave us lacking a certain type of understanding. Consider an extreme case of using non-portable properties, where we explain every different fact by appealing to a different non-portable property. For example, imagine we explain e by saying that the world has property P at t. And we explain f by saying that the world has property Q at t_1. And we explain g by saying that the world has property R at t_2, and so on. Not only are such explanations at odds with scientific practice they also seem unenlightening.

I don't have a full account of the nature of scientific understanding that I can appeal to to say why such explanations are unenlightening (though I will discuss this further later). But I suspect that the reason they seem unenlightening is because the collection of such explanations portrays the world as unstructured. The above explanations don't reveal the connections between e, f and g; we don't understand how these explananda and the corresponding explanations relate to each other. Regardless of this hypothesis, though, the key thought is just that there is something suboptimal about the explanations involving non-portable properties.

3.2 Connections to other Properties

The second way in which the property P seems gerrymandered is that it lacks substantial connections to other properties. To start with, notice that P is not interestingly connected to the properties involved in e, the fact about good G rising in price. The property P is not

[6] Strictly, the properties are things like *being an agent*, *having a certain degree of demand* and so on.

generally related to price rises, for example. Although P is closely connected one particular instance of prices rising – the world being P at t picks out all and only the worlds where e holds – it is not so connected to other instances.

More generally, some sets of properties have rich and substantial explanatory connections to each other — the Newtonian and microeconomic cases mentioned above illustrate this. Force, mass and velocity are deeply explanatorily connected to each other — facts about velocity are very often explained by facts about force and mass, for example. Similarly, the microeconomic properties are very closely explanatory connected in a variety of ways. Facts about demand for goods are explained by agents' preferences; facts about certain preferences are explained by other preferences; facts about certain choices are explained by preferences and prices; facts about prices are explained by facts about demand and supply; facts about the existence of certain goods are explained by the demand for other goods; facts about the supply of goods are explained by the demand for certain factors of production; facts about the demand for factors of production are explained by the price of the goods that they are used in producing; and so on.

The property P does not seem to have such rich and substantial connections to any such set of properties. There is a sense in which the property P seems 'isolated' from other properties.[7]

3.3 Explanatory Networks

These two considerations, of the portability of P and it's connection to other properties, together point towards an understanding of what is wrong with using P in explanations.

[7] You might worry whether P really is isolated in this way, surely there is some set of, perhaps unnatural, properties that P has substantial explanatory connections to. This point will be discussed at length in section 5.2.

When we give an explanation of a fact we don't want to appeal to a property, like P, that seems to be cooked up specifically and only for that explanatory purpose — such a property seems inappropriately gerrymandered. Rather, our explanatory practice favors appealing to properties that have a much richer set of explanatory connections – properties that we can use to explain a variety of facts, and which have close explanatory connections to other properties. Our practice, that is, favors properties that are embedded in *rich explanatory networks*.

We can see this by observing our scientific practice, but furthermore, as we noted in the discussion of portability, there is a thought about understanding here: we seem derive a certain understanding from theorizing in terms of properties that are deeply connected to each other and are part of a rich explanatory network.

Consider, for example, price changes in various goods. One good rose in price because of increased demand. Another increased in price because of decreased supply. A third good's price remained stable because while demand increased, supply also increased in a compensating way. A fourth good's price remained stable because supply and demand remained stable. A fifth good's price fell because although demand increased, supply also increased by a very large amount.

Seeing the connections between these facts, and how they are explained in related ways, greatly helps our understanding of those facts and of price changes more generally. (In fact, it is by illustrating the relations between such cases that introductory economics textbooks attempt to elicit understanding.) We can see the connections between these facts only because we, in explaining them, appealed to properties like demand, supply and price that are portable to various explanatory contexts and are closely explanatorily related to each other — that is, by appealing to properties that are part of a rich explanatory network.

Again, I don't have an account of understanding to give here, but previous discussions of understanding do fit with these thoughts. For example, Grimm (2012, p. 103), in his survey of the literature, characterizes the notion of understanding discussed by epistemologists as 'roughly, the good of being able to "grasp" or "see" how the various parts of the world were systematically related'. The explanations of the price rises in terms of demand and supply help give us this type of understanding, while explaining using non-portable properties does not allow us to see these systematic connections.

But even without any theoretically heavy notion of understanding, I think it is clear that elucidating these systematic connections between properties give us a kind of explanatory benefit with respect to a subject matter. Isolated explanations give less understanding than a unified body of explanations.

The core idea of this section is that the properties we want to use in theorizing — the natural properties — are part of rich explanatory networks: they are portable to various explanatory situations and they have substantial connections to other properties. On the other hand, gerrymandered properties like P don't have such features. In the next section I'll give an account of naturalness based on this idea.

One quick point before we move on to that though: An account of naturalness based on these ideas has the added advantage that it might allow us to answer a problem pressed recently by Dasgupta (2018), but which goes back further in the literature. That is, assuming that there are natural properties, why is it the case that we should theorize using such properties? Why should naturalness, as Dasgupta puts, it be *theory-guiding*? Of course, nothing said here helps explain why *F-naturalness* is theory-guiding. But if it's right that theorizing in terms of properties that are part of rich explanatory networks helps us gain understanding then an account of special science naturalness based on such networks goes a long way towards dealing with Dasgupta's challenge as it applies to the special science case.

4 The View

The guiding idea, from the last section, is that perfectly special science natural properties are part of rich explanatory networks. (Other properties have lesser grades of naturalness depending on their definition in terms of the perfectly natural properties. More on this later in the section.)

In this section I will develop this guiding idea, by giving an account of what makes for rich explanatory networks and how they connect to naturalness.

But just before that we should deal with an objection to this type of approach.

4.1 A Circularity Concern

The view I'm developing appeals to explanation to make sense of naturalness — natural properties are part of rich explanatory networks. But, there is a naturalness constraint on explanation — we discussed this in section 1.1. The existence of this constraint was part of the motivation for requiring an account of naturalness. But if naturalness is a prerequisite for explanation then giving a reductive account of naturalness in terms of explanation seems to be off the table.

I think we can deal with this problem and still give a reductive account of naturalness that fits with the idea that natural properties are part of rich explanatory networks. The move is to ignore, so to speak, the naturalness constraint on explanation when we are giving an account of naturalness. More specifically, consider explanation*. Explanation* is just the relation of explanation but with the constraint that only natural properties can be used to explain removed. An explanation* would also be an explanation if the properties used to do the explaining* were natural. We should give the account of naturalness in terms of explanation* not explanation. So, natural properties are members of rich explanatory* networks.

Of course, this idea of 'subtracting' the naturalness constraint from the account of explanation is easier to implement with some accounts of explanation than with others. In this paper, I'm going to be working with an account of explanatory correctness inspired by Lewis (1986a). The account is a very minimal one, it just says that A explains B when A gives information about the nexus of causal relations that led up to B and where A doesn't contain very unnatural properties and is not very disjunctive. Obviously, it is trivial to make this into an account of explanation* — just delete the clause about naturalness.

However, as I'll discuss more in section 4.3, although I'm working with this simple account of explanatory correctness my account isn't tied to it. I could appeal to other accounts of explanation too, as long as I remove the naturalness constraint. There are some accounts where we might need to do a little more work to factor out the component that appeals to special science naturalness, though. For example, certain causal accounts of explanation may build in special science naturalness by having a naturalness constraint on causation. In this case we would need to extract the naturalness constraint at the point of the account of causation.

Perhaps there are accounts of explanation where special science naturalness is so deeply enmeshed in the account that there is no way to extract it and be left with anything coherent. My account would not be able to appeal to those accounts of explanation to give a reductive account of special science naturalness. This, then, is a commitment of my approach.

But in any case, for the rest of the paper I'm going to work with the minimal causal account of explanation and the corresponding account of explanation*.

However, there is a concern about this move to explanation*: I motivated this approach to naturalness by appealing to the connection between naturalness and certain *explanatory* concepts. Consequently, the appeal to explanation* might seem like a change of subject that takes us away from these motivating ideas.

This is a reasonable concern, but it turns out there is no problem here. We can define naturalness in terms of explanation* but in a way that recovers the relevant biconditionals between naturalness and explanation. My account is going to say that properties are perfectly natural if they have certain explanatory* characteristics. But, if such properties are perfectly natural in virtue of having these explanatory* characteristics then they are going to have the analogous explanatory characteristics too, because the only difference between explanation* and explanation is whether it requires naturalness. So, an account which says that a property is perfectly natural if and only if it is part of a rich explanatory* network will also say that a property is perfectly natural if and only if it is part of a rich explanatory network.

4.2 Explanatory Clusters

We are now in a position to give the account of naturalness. I'll start by stating the account fairly directly — I'll then take some time to discuss and explain it.

Consider the concept of *explanatory* clustering*. Explanatory* clustering is a feature of a set of properties. It is a graded notion — in particular it is graded in three ways. The explanatory* clustering of a set of properties Γ increases with increases in:

(i) The proportion of the facts about properties in Γ that are explained* by other facts about properties in Γ

(ii) The goodness of these explanations*

(iii) The number of facts about properties in Γ

So, a set of properties scores very highly on explanatory* clustering if *most* of the facts about properties in Γ are explained* *well* by other facts about Γ and there are a *large number* of facts about Γ.

The basic idea then is that a property is perfectly natural if and only if it's part of sufficiently explanatorily* clustered set of properties. It turns out we need a couple of tweaks on this basic idea, so the account of natural properties is as follows:

A property is perfectly natural if and only if (i) it's part of sufficiently explanatorily* clustered set of properties, (ii) none of the properties in the set are definable in terms of the other properties in the set and (iii) that set is a *local maximum* of clustering — that is, we can't increase the explanatory* clustering by making small changes to the properties in the set.

The second condition is required because otherwise my account would imply that if A and B are perfectly natural then A&B is perfectly natural. The third condition is required otherwise we could take a set of properties that is well over the required threshold for explanatory* clustering and then add an arbitrary property and the resulting set would still likely be over the threshold. And consequently the arbitrary property would count as perfectly natural.

To simplify the terminology let's call a set of properties that meets conditions (i)-(iii) an *explanatory* cluster*. So, a property is perfectly natural if and only if it is part of an explanatory* cluster.

Notice, this implies that a property is perfectly natural if and only if it is part of an explanatory cluster (note the lack of an asterisk on 'explanatory'). This is because if an explanatory* cluster is made up of natural properties, then it is an explanatory cluster and if something is an explanatory cluster then it is also an explanatory* cluster because the conditions on explanatory* clustering are strictly weaker than those on explanatory clustering.

This notion of an explanatory cluster captures, and is motivated by, the ideas about explanatory networks that were developed in the last section. In particular, in the last section we discussed two characteristics which very unnatural, gerrymandered properties like P do not have, but natural properties do have. The first was to do with portability — the gerrymandered property P is not portable since it is not useful in explaining a wide range of explananda. The second was to do with connections to other properties — P does not have explanatory connections to other properties, rather it is explanatorily isolated. In these two senses, P is not part of an explanatory network while acceptable special science properties, like those of microeconomics, are.

If a set of properties is an explanatory cluster then the members will have the characteristics of being portable and having substantial connections to other properties. Imagine a set of properties $\{A, B, C, D\}$ that form an explanatory cluster. Then there are a large number of facts about those properties and most of them can be explained well in terms of facts about A, B, C or D. This means that A, B, C and D are portable — they are useful in explaining a wide range of facts — facts about A, B, C or D.

Similarly, the properties A, B, C and D are tightly explanatorily connected to each other. There are many explanatory connections between the facts about these properties. In fact, although this is not guaranteed by anything said so far, if large numbers of facts about A, B, C or D are explained by other facts about those properties, then it is likely that there are some systematic explanatory connections between these properties. That is, it is likely that there are fairly uniform, perhaps lawlike, ways in which certain combinations of instantiations of the properties A, B, C and D, explain other instantiations of the properties. If this were not the case then it would be extremely coincidental that most of the facts about those properties were explained by other facts about those properties.

So, sets of properties that are explanatory clusters form the rich explanatory networks that we discussed in the previous section.

This is an account of perfectly special science natural properties. Given this account we can define a graded notion of naturalness in roughly the way mentioned at the end of section 2.3. A property is very special-science-unnatural if it's definition terms of the perfectly special science natural properties is long, just as a property is very F-unnatural if its definition in terms of the perfectly F-natural properties is long. In the literature on graded F-naturalness, it is common to add that disjunctive definitions give rise to less natural properties that conjunctive ones (see, for example, Dorr and Hawthorne, 2013, pp. 16-17). I will say the same thing about graded special science naturalness.

(Note that we can't use the graded notion of clustering to give us an account of graded special science naturalness, because of the concern, mentioned earlier in the section, that we could take a set of properties scores very well explanatory* clustering and then add an arbitrary property and the resulting set would still score well on clustering, thus suggesting that the arbitrary property is fairly natural.)

That is the core of the account. In the rest of this section I'm going to go into more detail on a couple of parts of this account — the appeal to *explanatory goodness* and the appeal to *aboutness* — and I'll discuss some notable features of the account.

4.3 Explanatory Goodness

The account of special science naturalness appeals to explanatory goodness. (Or rather, strictly it appeals to explanatory* goodness, but I'll come back to this soon.) An account of explanatory goodness tells us when explanations are better to worse. What's to come will not really require us to fix on one particular account of explanatory goodness. Rather, a loose intuitive grip on explanations being better or worse is all that we need.

But for definiteness, and to help us get an intuitive grip on explanatory goodness, I will briefly outline an account of explanatory goodness that we might use. My preferred account is based on the idea that there should be a kind of proportionality between explanans and explanandum. The explanans shouldn't be far more specific than the explanans — it shouldn't be the case for example, that the explanans contains a huge amount of physical detail but the explanandum is a very general fact that doesn't imply such physical detail. But neither should it be the case that the explanans is a very general fact while the explanandum is very specific. This kind of idea is very much in the spirit of Yablo's (2002) approach to causation.

The account starts by assuming some account of explanation. It then builds explanatory goodness out of this. For simplicity, let's work with the very simple account of explanation that we mentioned earlier. The account says that A explains B when A gives information about the nexus of causal relations that led up to B. But, as we have discussed, there has to be a naturalness constraint too — the explanans cannot contain very unnatural properties or be very disjunctive.

Given this, we can outline two dimensions of explanatory goodness. The first dimension of explanatory goodness, PRECISION, says that explanations are better if in more of the physically possible worlds where the explanans is true, the explanans explains the explanandum.

For example, the explanation of an ice cube melting from the fact that is was dropped in warm water scores very highly, but not maximally, on PRECISION because nearly all the microstates that could realize the system are such that dropping the ice cube in warm water does lead to the melting. The explanation of the glass of water spilling because I knocked the table on which it was sitting with my knee scores worse on PRECISION, because there are many physically possible worlds where I knock the table with my knee but that does not lead to any spilling (for example, those where I don't knock the table hard enough).

The second dimension, ROBUSTNESS, says that explanations are better if in more of the physically possible worlds where the explanandum is true, the explanans explains the explanandum. An explanation of the extinction of a species which cites the introduction of a predator will score higher on ROBUSTNESS than the explanation of the same fact which cites the fundamental laws and the fact f which describes the total state of the world at a particular time t. That is because more of the physically possible worlds in which the species goes extinct are such that the extinction is explained by the predator's introduction than explained by f. f cannot explain the extinction in many worlds because it holds in only very few physically possible worlds.

ROBUSTNESS tells us that that the explanans shouldn't hold in few possible worlds when the explanandum holds in many. And PRECISION tells us that the explanans shouldn't hold in many worlds when the explanandum holds in few. Together they tell us that there should be a proportionality between explanans and explanandum.

Of course, there is a lot more to say about this approach to explanatory goodness. I discuss it at length elsewhere. But this is just an example of one type of approach to explanatory goodness that could fit into the account of naturalness I'm developing here. But again, the discussion of naturalness going forward is not deeply tied to this particular approach to explanation.

However, strictly speaking my account of naturalness appeals to explanatory* goodness — a sense of goodness that applies to explanations* not explanations. But it is easy to adapt the account of explanatory goodness we just outlined into one of explanatory* goodness. Just remove the naturalness constraint from explanation, thus yielding an account of explanation*, and the same definitions of PRECISION and ROBUSTNESS and give us an account on explanatory* goodness.

4.4 Aboutness

As well as appealing to explanatory* goodness the account of clustering also appeals to the idea of a fact being *about* a property. The issue of 'aboutness' is extremely complicated. It is not at all clear what it is for a fact to be about — or to be statable in terms of, or to have as part of its content — a particular property or individual. There are many difficult cases. For example, consider logical truths that contain non-logical vocabulary, like 'everyone is either a philosopher or is not a philosopher'? Is that about the property of *being a philosopher*? Or consider universal generalizations. Is a universal generalization in part about a particular individual that is an instance of the generalization? Is, for example, the generalization 'all philosophers are human' partially about me, a human philosopher?

I'm not going to go into these difficult cases, and I'm certainly not going to give a theory of aboutness. Rather, I'm going to mainly rely on some intuitive judgements regarding aboutness.

Here are a few (undoubtedly exception-ridden) generalizations that we can make. Consider a set of properties Γ. When we have a fact that ascribes one of those properties, or a combination of such properties, to an individual then that fact is about (some of) the properties in Γ. So, to go back to our Newtonian set of properties {mass, position, force}, a fact that ascribes a particular mass to an object is about a property in Γ. Sometimes a fact ascribes a property to an individual that is defined in terms of the properties in Γ in this case again, the fact is about the properties in Γ. For example, a fact that ascribes a particular momentum to an individual is also about the Newtonian set of properties given how momentum can be defined in terms of those properties. If we have a fact that ascribes a property to an individual that cannot be defined in terms of Γ then that fact is not about the properties in Γ.

Plausibly there are exceptions to these claims. But in most cases our intuitive judgments of when a certain fact is about a property are relatively clear. Even if there are some cases where it is not clear whether some fact is about some property this does not matter too much for my ends. The appeal to aboutness in the account of clustering does not rely on fine judgements of aboutness in specific cases, rather the account only requires much coarser judgements regarding 'most of the facts about Γ'. In particular, we can often be confident that most of the facts about Γ are explained by other facts about Γ without having sharp views on aboutness in every case.

4.5 Some Structural Features of the Account

Now that we have the account on the table there are a few important features of the view that we should note.

(1) We can see what determines the facts about special science naturalness. The account of naturalness is given in terms of explanation* and explanatory* goodness. And if we assume the account of explanatory* goodness I suggested above, then the facts about special science naturalness are ultimately grounded in causal facts.

(2) Since naturalness is grounded in explanatory* facts, then facts about special science naturalness are *contingent*. This contingency is because naturalness of a property, on my account, isn't some deep metaphysical fact, rather it is tied to how science inquires into the actual world.

(3) The account is holistic — the naturalness of a property is grounded in it's connection to other properties. A property isn't natural in virtue of some special intrinsic feature it has. Rather, it is natural because it is closely explanatorily tied to cluster of other properties.

(4) These features suggest that my account should be amenable to people who are suspicious of the metaphysical profligacy of traditional F-naturalness talk. The account appeals to facts that the metaphysically scrupulous should be comfortable with.

(5) It's an advantage of my view, though not one that I'll be able to discuss in great detail here, that explanatory clusters are sets of properties which particularly invite autonomous scientific inquiry. If a set of properties Γ is an explanatory cluster then most of the facts about Γ are well explained by other facts about Γ. We do not need to look outside Γ, so to speak, in order to give good explanations. Γ, therefore, defines a subject matter that is particularly amenable to autonomous investigation. So, the account of naturalness goes a long way towards explaining the way in which special sciences often progress relatively autonomously of each other and of physics.

5 Explanatory Clusters

Now that we have this account, we need to consider it gets the right results. Does the account line up with our judgments about naturalness in the practice of the special sciences? Of course, answering this question fully is a huge project, it involves looking in detail at various special sciences and considering the various causal and explanatory relations in order to see how the putative natural properties cluster together. Right now, I'm just going to outline some general considerations which should lead to optimism that my account gets the results that we want in the majority of cases, and some other considerations that suggest that my view has limitations.

Let's take this in two steps. Firstly consider whether the account allows in the properties we want; that is, whether it classifies the properties used in paradigm special science explanations as natural. After that, consider whether the account rules out unnatural properties — like the property P.

5.1 Letting in Natural Properties

Consider microeconomics and a set of properties that are fairly distinctive to that field, for example, *agent*, *goods*, *preference*, *demand*, *supply*, *factor of production*, *choice*, *price*.[8]

It looks like this set will count as highly explanatorily clustered because, as we noted in section 3, these properties are explanatorily interconnected in a large range of ways. Facts about demand for goods are explained by agents' preferences; facts about certain preferences are explained by other preferences; facts about prices are explained by facts about demand and supply and so on. In fact, standard microeconomic theorizing tells us about these connections in very systematic ways; there is a large body of theory regarding exactly how a few facts about these properties explain the others. It looks like, then, these microeconomic properties are part of an explanatory cluster and so count as natural.

For another example consider classical genetics. *Gene*, *allele*, *trait*, *dominance* and *inheritance*

[8]This is not to say that this is the unique set of basic microeconomic properties. Rather, this is just a list of a few properties that are particularly central to microeconomic theorizing.

are all closely connected by good explanations and will form a cluster. Further, consider population ecology and properties like *population*, *generation*, *predator*, *prey*, *carrying capacity*; and thermodynamics and properties like *temperature*, *pressure*, *entropy*.

But we can do more than just survey examples of special sciences. We have more general reasons to think that this account of naturalness will allow in paradigm special science properties – the ones that are part of good explanatory theories of particular special sciences. A theory of a particular special science will be framed in terms of a set of properties Γ with axioms about those properties and rules for inferring other propositions about those properties. If the theory is successful it will be able to explain well many facts about Γ. In this case, the properties of Γ would form an explanatory cluster because some facts about Γ would explain well a large number of other facts about Γ.

When we have special sciences, like microeconomics and classical genetics, for which we have good explanatory theories then the basic properties of those theories will likely form an explanatory cluster. But there are special sciences for which we do not possess such a successful explanatory theory. However, often in such cases we may have good reason to believe that a theory could be developed in the future. If this is so then we have some reason to think that that the properties that the special science is currently using are natural. Perhaps neuroscience and macroeconomics are fields where our current theorizing has progressed far enough that it is reasonable to be optimistic in this way.

But there may well be cases of current special sciences where there isn't good reason to think that such an explanatory theory can be developed. In this case, I think we have little reason to think that the basic properties of those special sciences form explanatory clusters, and so little reason to think that they are natural. But this is, I think, an acceptable result. The story about explanatory clusters tells us what natural properties the world presents us with.

But there's no reason to think that it's easy for us to find out what these natural properties are, and so no reason to think that current scientific practice always maps on to the natural properties in the world.

5.2 Ruling out Unnatural Properties

So the account gets reasonable results in classifying properties actually at work in the special sciences as natural. Does the account rule out the right properties? Or can we have explanatory* clusters that contain intuitively unnatural properties? (I'll call these 'unnatural clusters' for short.)

Let's start by seeing why it is not easy to construct unnatural clusters. The key idea is that my account of explanatory* clustering, and hence my account of naturalness, involves a certain kind of 'closure' that makes it hard to cook up unnatural properties that meet my criteria of naturalness. To see this, let's look at the example of the property P again.

As a reminder, the explanatory* clustering of a set of properties Γ increases with increases in:

(i) The proportion of the facts about properties in Γ that are explained* by other facts about properties in Γ
(ii) The goodness of these explanations*
(iii) The number of facts about properties in Γ

Imagine that we try to add a cooked up unnatural property, like P, to some other set of natural-looking properties. P was designed to give a good explanation* of a particular microeconomic fact e. The fact that the world had property P at t gives a good explanation* of the fact e about the price of a particular good rising. So what would happen if we added

P to the the standard set of microeconomic properties?

As we noted above, there is a set of microeconomic properties, call them Δ, that score well on explanatory* clustering. However, adding P to this set of properties would not increase the clustering — that is $\{\Delta \cup P\}$ is not more clustered than Δ.

It is true that facts about P — in particular, the fact that the world had property P at t — would explain* well one particular fact about Δ, the fact e. But this is not enough for $\{\Delta \cup P\}$ to be more clustered than Δ. This is because the facts about Δ and P together do not explain* well *any of the other* facts about P. There are a large number of facts, actually most of the facts, about P, that are not well explained* by the properties in $\{\Delta \cup P\}$. For example, the the fact that the world does not possess property P at time t_4 is not well explained* by the properties in $\{\Delta \cup P\}$. So $\{\Delta \cup P\}$ scores worse than Δ on clustering. And consequently $\{\Delta \cup P\}$ cannot be an explanatory* cluster because it is not a local maximum of clustering.

This idea generalizes to other gerrymandered properties. Gerrymandered properties like P will not generally increase clustering when combined with traditional special science properties because the special science properties won't explain well most of the facts about P. For a property to count as natural it really needs to be deeply integrated with a cluster of other properties — it's not enough to just have the relatively minimal connections that P does to the microeconomic properties.

But what if we manage to cook up a property Q that does fit closely with the traditional economic properties — one that explains* well lots of the facts about the traditional economic properties and where many of the traditional economic properties explain* well facts about Q? Well in this case it's reasonable to accept that we have just discovered a new natural economic property. If there really are these rich and deep connections between Q and the other economic properties then Q should be integrated into our economic theorizing in the way that other natural economic properties are.

And this is the case even if Q seems unnatural to us on first glance. Many special science properties can seem unnatural when first introduced, and only come to seem natural once we are familiar with the theories in which they are embedded. For example, considered in itself, *gene* might seem fairly unnatural — it is by seeing how genes are related, for example, to traits organisms possess and to inheritance of those traits that we come to find the concept natural.

But maybe the way to cook up an unnatural property that meets my criteria isn't to add a property to some traditional special science properties but to build an explanatory* cluster of properties, all of which are intutively unnatural, *from scratch*. I don't think it's particularly easy to see how to do this — there doesn't seem to be any simple recipe for constructing clusters. But I'm sure it is possible.

However, I don't take this to be a problem for the view. I think its a reasonable result that there are such explanatory* clusters and members of such clusters are natural. There is no reason to think that we have discovered all the special science natural properties that are out there, and no reason to think that we have discovered all the special sciences. If we find a novel set of properties that form a rich explanatory* network, where there are lots of facts about those properties and where most of the facts about those properties explain* well other facts about those properties then we have found a subject matter that looks similar to paradigm special sciences. So, we should be open to accepting that those properties are natural and that we have discovered a new special science.

5.3 A Limitation of the Account

So, the strategies of adding a gerrymandered property to a traditional special science properties, and that of constructing a new cluster of unnatural properties from scratch don't cause a problem for the account so far. However, as we have foreshadowed, there is a limitation to the account. That limitation is due to another strategy for constructing clusters containing intuitively unnatural properties. This strategy involves taking a whole explanatory cluster and then permuting all the properties in it in ways that retain the relations between the properties that lead to clustering. Consequently, the account of naturalness doesn't rule out every intuitively unnatural property.

In particular, clustering is to do with facts about a set of properties explaining* and explaining* well other facts about those properties. The account of explanation* we have been working with is a simple causal account: A explains* B when A gives information about the nexus of causal relations that led up to B. And the account of explanatory* goodness we've been working with is to do with explanations* holding across physically possible worlds.

So, if we can take an explanatory* cluster and transform each property in it such that the facts about those transformed properties have the same causal relations as the original properties, and these causal relations hold in the same way across physically possible worlds as the original properties, then this transformed set will also be an explanatory* cluster.

And there are some ways in which we can perform such a transformation. The simplest is a kind of 'property switching' — of the kind exhibited by the classic definition of *grue*. Let grue be green before the year 2000 and blue after. And let bleen be blue before 2000 and green after. One intuitive way to describe this is as grue and bleen 'switching roles' in the year 2000, previously grue played that green-role, after it plays the blue-role, and similarly with bleen.

Of course, this idea can be extended to more than just two properties. So if we have an explanatory* cluster of properties, P, Q, R, S and T, we can form a transformed set of properties, A, B, C, D and E where each of the transformed set 'switches roles' at a specific time. So, for example, A plays the P-role up to 2000 and the Q role after, and so on for the other transformed properties. It looks like the transformed set of properties will also be an explanatory* cluster. This is because the causal relations between facts about the transformed set will mirror the causal relations between facts about the original set of properties, in both this and other physically possible worlds. (There are a lot of complexities about how exactly such transformations will work and in what cases the causal relations are mirrored in the transformed set. But discussing these details would take us too far astray.)

Other transformations might work too. For example, perhaps 'time-shifting' every property in a cluster — for example, replacing the property P with the property *ten-minutes-before-P*, and so on — can form a new cluster. There might be other possible transformations too.

One way to put to put it, then, is that for every explanatory cluster there is an equivalence class of such clusters joined by transformations like the ones described above. So, the account of naturalness puts you in the right equivalence class, so to speak, but it does not go further to pick out the right cluster in each equivalence class. So although the account does rule out the gerrymandered properties discussed in section 3, it doesn't rule out a special set of unnatural properties: the grue-ish properties produced by property-switching transformations and similar transformations.

This is a gap to be plugged in order to get full necessary and sufficient conditions on naturalness. It would be desirable to add something to the account here to pick out the right cluster from each equivalence class. I'm undecided on the best way to do this.

Perhaps the most obvious option is to appeal to pragmatic considerations. Out of each equivalence class of clusters we could, for example, pick out one that contains the properties that are most psychologically salient to us, or which are most closely connected to ones that are psychologically salient to us. Or we could appeal to other pragmatic considerations like which properties are most easily manipulable by us.

Alternatively, we could try to adapt other approaches to special science naturalness and put them to work in this restricted context of just picking out one cluster from each equivalence class. Though I'm not especially convinced that the problems with, for example, the graded F-naturalness approach, are any easier to deal with in this context.

Or, perhaps, there is some additional virtue that clusters might have that we can appeal to – perhaps to do with the simplicity with which the explanatory relations between the properties in the cluster could be expressed. But it's not easy to see how to develop this to get us the results we want.

So, I think picking out the right cluster from each equivalence class is a real challenge. Consequently, the account developed here does not get us everything that we want. It doesn't give us a full set of necessary and sufficient conditions for special science naturalness.

But it does do a lot — it gives us a picture of what naturalness could be in the special sciences that is deeply connected to the practice of science, that makes sense of why we might care about using natural properties in our theorizing and which successfully rules out the types of gerrymandered properties we discussed in section 3 while allowing in paradigm special sciences properties of the type discussed in section 5.1.

So while there's more work to be done on how to fill out this account, I think the account of this paper is a very promising place to start.

And there's more work to be done along another dimension too. My focus here has been on the connection of natural properties to explanation. But, as we noted in section 1, naturalness has connections to more than just explanation. So the project would continue by considering which of the putative roles of natural properties could be played by natural properties in my sense. Given the close ties between explanation and, for example, laws and causation, it seems plausible that the properties I've identified can play these other roles of natural properties. But that work will have to be done elsewhere.

References

Armstrong, D. (1983). *What is a Law of Nature?* CUP.

Butler, D. and B. E. Cain (1992). *Congressional redistricting: Comparative and theoretical perspectives*. Prentice Hall.

Clapp, L. (2001). Disjunctive properties. *Journal of Philosophy 98*(3), 111–136.

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

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

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

Fodor, J. A. (1974). Special Sciences. *Synthese 28*, 97–115.

Franklin-Hall, L. R. (2016). High-level explanation and the interventionist's 'variables problem. *British Journal for the Philosophy of Science 67*(2), 553–577.

Grimm, S. R. (2012). The value of understanding. *Philosophy Compass 7*(2), 103–117.

Hawthorne, J. (2007). Craziness and Metasemantics. *Philosophical Review 116*, 427–440.

Kitcher, P. (1981). Explanatory unification. *Philosophy of Science 48*(4), 507–531.

Kitcher, P. (2001). *Science, Truth, and Democracy*. Oxford University Press.

Lewis, D. (1983). New work for a theory of universals. *Australasian Journal of Philosophy 61*(4), 343–377.

Lewis, D. (1986a). Causal explanation. In D. Lewis (Ed.), *Philosophical Papers Vol. Ii*, pp. 214–240. Oxford University Press.

Lewis, D. (1986b). *On the Plurality of Worlds*. London: Blackwell.

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

Loewer, B. (2009). Why is there anything except physics? *Synthese 170*(2), 217–233.

Sider, T. (1995). Sparseness, immanence, and naturalness. *Nous 29*(3), 360–377.

Sider, T. (2011). *Writing the Book of the World*. Oxford University Press.

Strevens, M. (2008). *Depth: An Account of Scientific Explanation*. Harvard University Press.

Weslake, B. (2015). The problem of disjunctive explanations. *manuscript*.

Yablo, S. (2002). Mental Causation. *Philosophical Review 101*, 245–280.