

A Defense of Objectivism about Evidential Support

Abstract

Objectivism about evidential support is the thesis that facts about the degree to which a body of evidence supports a hypothesis are objective rather than depending on subjective factors like one's own language or epistemic values. Objectivism about evidential support is key to defending a synchronic, time-slice-centric conception of epistemic rationality, on which what you ought to believe at a time depends only on what evidence you have at that time, and not on how you were at previous times. Here, I defend a version of objectivism about evidential support on which facts about evidential support are grounded in facts about explanatoriness.

Keywords: Diachronic norms, epistemology, rationality, evidence, explanation

1 Introduction

My aim in this paper is to defend the claim that evidential support is an objective matter against a challenge from subjectivists. Subjectivists about evidential support hold that a body of evidence only supports a given hypothesis relative to some subjective factor, like a set of evidential standards, or standards for evaluating evidence, and that there is no privileged set of evidential standards. For subjectivists, the only case where there is a non-relative fact of the matter about the degree to which evidence supports a hypothesis is the trival one where the evidence entails the hypothesis or its negation, for this is a case where all permissible evidential standards agree on the degree to which the evidence supports the hypothesis. By contrast, objectivists hold that whether and to what extent a body of evidence supports a hypothesis is an objective matter, so that a body of evidence supports a hypothesis to a given degree *simpliciter*, not just relative

to a choice of evidential standards. (Objectivists can endorse talk of evidential standards; they'll just hold that there is a single privileged choice of evidential standards, whereas subjectivists think there are no such privileged standards.) Objectivists think of evidential support as a two-place function, taking as arguments a body of evidence and a hypothesis (and outputting a degree of evidential support), while subjectivists think that evidential support must be a three place function, taking as arguments a body of evidence, a hypothesis, and a set of evidential standards.

The debate between objectivism and subjectivism about evidential support is closely related to, but not identical with, another debate in epistemology—the debate over Uniqueness and Permissivism. Uniqueness theorists hold that there is a unique doxastic state that it is rational for one to be in, given a body of total evidence. Permissivists hold that for at least one body of total evidence, there is more than one doxastic state that it is rational for one to be in.

These two debates—between objectivism and subjectivism about evidential support, and between Uniqueness and Permissivism—are in turn central to the debate over diachronic norms on belief. Here's why. First, objectivism about evidential support is plausibly a necessary condition for Uniqueness to be true. If subjectivism were true, and so there were different equally permissible evidential standards, then it would be difficult to resist the thought that different agents could use different such evidential standards in arriving at their beliefs and still qualify as perfectly rational. Now, objectivism isn't quite a sufficient condition for Uniqueness. One might deny (one version of) the evidentialist thesis, which says that the only reasons for belief are evidential reasons. For instance, if there were pragmatic reasons for belief, then which doxastic states are perfectly rational would not be solely determined by facts about one's evidence (Kopec (ms) takes something like this approach, treating epistemic rationality as a species of instrumental rationality and hence sensitive to practical considerations). One might

also accept evidentialism but adopt a comparatively undemanding view of rationality. Perhaps the facts about evidential support are extremely precise, while agents cannot be reasonably expected to have such precise degrees of belief, either because they cannot know such precise facts about evidential support, or because they cannot achieve such fine-grained doxastic attitudes.¹ Still, objectivism about evidential support seems necessary for Uniqueness. If objectivism were false, it is difficult to see how Uniqueness could nonetheless be true.

Second, Uniqueness is plausibly a necessary condition for denying the existence of diachronic norms governing belief (and doxastic states more generally). Suppose we take it as a datum that it is irrational for an agent to have wild fluctuations in her beliefs (at least, unless her evidence likewise fluctuates wildly). To take one example, if she had the same evidence throughout the day but vastly different (*de dicto*) beliefs at different times in the day, this would be irrational.² The Uniqueness theorist can account for this datum without appealing to diachronic norms. Let the Uniqueness theorist espouse just the synchronic norm stating that at each time, your degree of

¹Note, however, that Uniqueness theorists may also simply deny that the facts about evidential support are so fine-grained. They might hold that the degree to which a body of evidence supports a hypothesis is sometimes to be represented by an interval of real numbers rather than by a single real number, for instance. Uniqueness theorists might also hold that it is sometimes indeterminate the degree to which a body of evidence supports a hypothesis, so that it is indeterminate what degree of belief an agent with that evidence should have in that hypothesis.

²Of course, the opponent of diachronic norms could also deny the alleged datum. Perhaps, for instance, it is permissible to change one's beliefs without a change in one's evidence if this change is driven by the adoption of a different set of epistemic standards. For instance, it might be permissible to come to give greater weight to simplicity in evaluating evidence, and this will lead one to revise one's beliefs without gaining or losing any evidence. Moss (2014) takes something analogous to this approach in discussing imprecise credences. She holds that an agent with imprecise credences (where her doxastic state is represented by a set of probability functions rather than by a single probability function) must identify with one member of that set of probability functions, and use that member as the basis for practical decision-making. She holds that it is permissible for that agent to come to identify with a different member of that set, where such a change is analogous to a rationally permissible change in her (practical) values. While Moss is discussing a change in which aspect of one's doxastic state one uses for practical decision-making, rather than a change in one's doxastic state itself (since she is assuming that one's overall imprecise credal state remains constant absent changes in evidence), one could go further and say that it is also permissible to undergo a change in one's doxastic state without a change in evidence due to a similar change of heart. Christopher Meacham (personal communication) tells me that he is sympathetic to this sort of position.

belief in a proposition ought to be the degree to which that proposition is supported by your evidence at that time. If the agent has the same body of total evidence throughout the day but different doxastic states at different points, then at least most of the time, her doxastic state was not the uniquely rational one, given her evidence. But this explanation is unavailable to the Permissivist, for she denies that there is always a uniquely rational doxastic state to be in, given a body of evidence. Instead, she must appeal to diachronic norms to proscribe wild fluctuations in belief. Suppose that each of the doxastic states the agent has during the day is individually permissible, given her total evidence. For instance, each is based on a different, but permissible, set of evidential standards. The Permissivist might then appeal to a diachronic norm stating that it is a rational requirement that one always base one's doxastic state on the same set of evidential standards (and that at any given time, one ought to have the doxastic state which is proportioned to her evidence, given that set of evidential standards). Different agents could rationally have different doxastic states in response to the same body of total evidence, but a single agent at different times could not.³

This point can be made precise in a Bayesian framework. We can think of a set of evidential standards as being represented by a prior probability function (or a *prior*, for short). Objectivists about evidential support think that there is a privileged prior which represents the privileged evidential standards. Subjectivists think that there are multiple permissible priors, since there are multiple permissible sets of evidential standards. Extreme subjectivists will hold that all probabilistically coherent priors are permissible, while more moderate subjectivists will hold that the set of permissible priors excludes some probabilistically coherent priors (the crazy ones), but is nonetheless not a singleton.

³Note that on this sort of Permissivist picture, epistemic rationality isn't all that much less demanding than on a Uniqueness picture. After all, maintaining the exact same evidential standards over time and always having the beliefs that those standards recommend is no mean feat!

Our Uniqueness theorist, who is also an objectivist about evidential support, can give the following theory about rational belief, which includes only synchronic norms: Let P be the privileged prior. If at some time an agent has total evidence E , then she ought to be in the doxastic state represented by the probability function $P(- | E)$. This theory will proscribe wild fluctuations in belief that aren't due to wild fluctuations in evidence. If an agent's probability function changes dramatically, and this isn't due to a dramatic change in her evidence, it must be due to her sometimes not using the privileged prior P .

Our Permissivist, however, will give a theory which includes some diachronic norm. Conditionalization is the most popular one. It says, in effect, that once one adopts a prior as one's own, one must stick with it. More precisely, it says that if at time t_0 one has probability function P_0 , and at t_1 one gains evidence E (and nothing stronger), then at t_1 one's probability function should be $P_1(-) = P_0(- | E)$.⁴

A small caveat: Uniqueness and Permissivism, as well as objectivism and subjectivism about evidential support, can be located on a spectrum. Permissivists can vary in just much wide they take the range of permissible doxastic state to be in any given evidential situation. For instance, very moderate Permissivists might think that it is permissible to assign some proposition a probability anywhere between 0.6 and 0.62, while more extreme Permissivists might hold that the permissible range is significantly wider, with the most extreme ones holding that any probability assignment between 0 and 1 is permissible. An analogous point can be made about subjectivism about evidential support. Moderate subjectivists will hold that there are some objective facts about evidential support that make it determinately the case that the evidence supports a given hypothesis to a degree that is somewhere in a fairly narrow range, but that only subjective factors, like the exact weights that one gives to various epistemic values (e.g.

⁴See Meacham's contribution to this volume for detailed discussion of subtleties involved in interpreting Conditionalization.

simplicity) could narrow things down further. More extreme subjectivists, by contrast, will hold that objective factors rule out far less. Whether the opponent of diachronic norms could settle for a very moderate form of Permissivism, and hence perhaps also a very moderate form of subjectivism, depends on why changes in one's doxastic state without changes in one's evidence are supposed to be irrational. If it is just that certain wild changes in one's doxastic state are intuitively irrational, the opponent of diachronic norms could perhaps be happy with a very moderate Permissivism. After all, such a moderate Permissivism would, even absent diachronic norms, allow only very minor fluctuations in one's doxastic state, and these aren't obviously intuitively irrational. If, by contrast, the theorist wants to prohibit non-evidence-driven changes in doxastic state on the grounds that such changes leave one vulnerable to exploitation along the lines of the diachronic Dutch Book for Conditionalization (Teller 1973, Lewis 1999), then the theorist must reject even very moderate Permissivism, since even very small fluctuations in one's doxastic state, without changes in evidence, will leave one in principle vulnerable to exploitation. Either way, the further along the spectrum toward full objectivism, the better for the opponent of diachronic norms. Going forward, then, I will be focusing on full objectivism, but it is important to keep in mind that the opponent of diachronic norms might be able to settle for something short of these extreme position.⁵

⁵In earlier work (Hedden 2015b, 2015c), I have defended a time-slice-centric picture of rationality on which all requirements of rationality are synchronic and impersonal. Obviously, this involves commitment to more than just Uniqueness, and hence objectivism about evidential support. For even if there are no diachronic norms governing what you ought to believe, there might be diachronic norms governing, say, the temporally extended process of reasoning, or evidence-gathering, or doxastic (as opposed to propositional) justification. Hlobil (2015) and Podgorski (forthcoming) argue that there may be diachronic requirements of rationality such as requirements governing the temporally extended process of reasoning, even if there are no diachronic norms governing what you ought to believe at a particular time. And even if there are no diachronic norms governing belief, there might be diachronic norms governing preferences, intentions, or actions, for instance. And even if there are no diachronic norms whatsoever, there might be norms that nevertheless fail to be impersonal, such as synchronic norms saying that you should defer to the attitudes that you *now believe* your early or later selves had or will have. Still, if objectivism about evidential support, or at least something close to it, is false, the whole edifice comes crashing down.

My aim in this paper is to defend objective evidential support against a certain sort of objection, pressed most forcefully by Titelbaum (2010). According to this objection, the fundamental facts about evidential support—what the privileged evidential standards are—must be *a priori* in order for evidential support to play the role in epistemology that objectivists want it to play, but the only plausible candidate factors that might determine the privileged evidential standards are *a posteriori* ones. (Throughout this paper, I use ‘determines’ in a metaphysical, rather than an epistemic sense, to mean something like ‘grounds’ rather than something like ‘discovers.’) Therefore, objective evidential support is either non-existent or cannot play the role that many of its advocates want (myself included), namely helping to determine what agents ought to believe. I will concede that it must be an *a priori* matter what the evidential support relation is, but I will argue that the most plausible factors that might determine this evidential support relation are not *a posteriori* ones after all.

2 Evidential Support, A Priority, and Language Dependence

The fundamental facts about evidential support must be *a priori*. The qualifier ‘fundamental’ is important here. Of course, it is often an empirical, *a posteriori* whether and to what extent some particular piece of evidence supports a given hypothesis. It is an empirical matter whether the evidence that there are dark clouds on the horizon supports the hypothesis that it will rain soon. For this depends on what other background evidence one possesses. This is compatible with the claim that it is an *a priori* matter whether and to what extent a given body of *total* evidence supports some hypothesis. Equivalently, it is compatible with the claim that it is an *a priori* matter what the objective evidential support relation is, or what the privileged evidential standards are.

In my view, these fundamental facts about evidential support must be *a priori*.

Before saying why that is, let me say how I am understanding the *a priori*. Following Field's (2000) characterization, some proposition is *a priori* if one has justification for believing it without empirical evidence, and this justification is not defeasible on empirical grounds. Similarly, a belief-forming method, like a set of evidential standards, counts as *a priori* just in case one has justification for using it that doesn't depend on empirical evidence supporting the method (call this 'default' justification), and moreover this justification for using it is not defeasible on empirical grounds. (If you dislike this characterization of the *a priori*, just replace '*a priori*' with 'Field-*a priori*' in what follows. For it is his notion of *a priority* that I think the evidential support relation must, and does, satisfy.)

Why, then, must the objective evidential support relation (if such exists) be *a priori*? I think there are two reasons, corresponding to each of Field's two conditions. First, the *a priority* of the objective evidential support relation may be necessary in order to avoid a pernicious epistemic circularity. If what the evidential support relation is depends on some fact F, and if it's an empirical matter whether F is true, then in order to proportion one's beliefs to one's evidence, one has to know whether F is true, but one would already have to know whether F is true in order to proportion one's beliefs to one's evidence (Titelbaum 2010). So, the idea is that in order for facts about evidential support to be able to play the role of determining what one ought to believe, they must be knowable *a priori*. Now, I am somewhat skeptical that one actually has to know what the evidential support relation is in order to (justifiably) proportion one's beliefs to one's evidence. Justified belief at the first order is compatible with lack of knowledge or justified belief at higher levels; one may justifiably believe H without knowing or even justifiably believing that H is supported by one's evidence, and certainly without knowing all the other facts about evidential support. Similarly, one might justifiably

make various deductive inferences without knowing what rules one is following and without knowing that they are valid inference rules.

Still, it is plausible that what one ought to believe supervenes on one's mental states.⁶ (This is a characterization of what Conee and Feldman (2001) call 'mentalist internalism,' which to my mind is more plausible than 'access internalism,' according to which what one ought to believe depends on factors to which one has some sort of special access.) If evidential support determines what one ought to believe, this means that what evidential standards one ought to employ in proportioning one's beliefs to one's evidence should supervene on one's mental states. The most obvious (though not the only) way for this to happen is for it to be necessary what the privileged evidential standards are. Of course, this isn't quite to say that they must be *a priori*; in principle they could be necessary *a posteriori*. But it would be worrying for defenders of objective evidential support if we didn't at least have default justification for employing these necessarily privileged evidential standards in forming one's beliefs. This corresponds to the first clause in Field's characterization of the *a priori*.

Second, it is a desideratum on any set of belief-forming rules, or any 'inductive method,' that it treat its most basic principles as empirically indefeasible. This corresponds to Field's second clause. The argument comes from Lewis (1971) (see also Elga 2010). Say that an inductive method is *immodest* if it always recommends itself over any competitors. It never recommends that, in response to a given piece of evidence, you adopt some other inductive method instead. Lewis cashes this out in terms of estimated accuracy: an inductive method is immodest just in case, given any body of total evidence, the inductive method's estimate of the accuracy of believing in accordance with that inductive method is at least as great as its estimate of the accuracy of believing in accordance with any competitor inductive method. Lewis (1971, 56)

⁶Indeed, this supervenience consideration is one of my main arguments for a time-slice-centric picture of rationality in Hedden (2015b; 2015c).

argues that any inductive method that is not immodest would thereby fail to deserve your trust:

Suppose you did trust some non-immodest method. By definition, it estimates some competing method to be more accurate than itself. So if you really did trust your original method, you should take its advice and transfer your trust to one of the competing methods it recommends. It is as if *Consumer Bulletin* were to advise you that *Consumer Reports* was a best buy whereas *Consumer Bulletin* itself was not acceptable; you could not possibly trust *Consumer Bulletin* completely thereafter.

An immodest inductive method treats itself as empirically indefeasible. So if there's an inductive method such that we have default justification for employing it, then while it can treat its non-basic methodological principles as empirically defeasible, it must treat its most basic such principles as not subject to revision on empirical grounds. If evidential support is to play the role of determining what one ought to believe (so that the relevant inductive method is to believe a hypothesis to the extent that it is supported by one's evidence), then what the privileged evidential standards are must not be subject to revision on empirical grounds.

So, objectivism about evidential support is committed to the claim that the privileged evidential standards are *a priori* in the sense that we have default justification for employing them in proportioning our beliefs to our evidence, and that the privileged evidential standards are not revisable on empirical grounds. Can an evidential support relation meet this requirement? What factors could determine the evidential support relation (or, equivalently, the privileged evidential standards), subject to this constraint?

It is useful to divide candidate facts into two broad types: formal (or syntactic) factors, on the one hand, and substantive factors on the other. Substantive factors are just non-syntactic ones. We'll get into some specific proposals later, but possible substan-

tive factors might include simplicity (if this isn't cashed out formally), explanatoriness, naturalness, beauty, and the like.

Carnap (1950) famously sought a theory on which substantive factors played no role, and instead formal, purely syntactic, factors would be sufficient to determine the degree to which a hypothesis is supported by some body of evidence. By espousing only formal factors, Carnap's theory would qualify as an inductive *logic* to stand alongside deductive logic. After all, deductive logic is often taken to be the theory of logical consequence. That is, it is the study of which sentences are consequences of which others purely in virtue of their form or syntactic structure. A sentence that follows from another partly in virtue of the meanings of its non-logical terms would not count as a *logical* consequence of it. 'All bachelors are messy' is a consequence of 'All unmarried men are messy,' but this is so partly in virtue of the meanings of 'bachelors' and 'unmarried men.' If we swapped out one of these predicates for another with a different meaning, the consequence relation might no longer hold. So what Carnap was seeking was a characterization of the confirmation or evidential support relation which was sensitive only to the syntactic structures, and not the meanings, of the sentences involved. If one sentence supports another to a given degree, then so should the result of swapping out all occurrences of some predicate 'F' in the two sentences with a new predicate 'F*' which doesn't occur in the original sentences.

It would be great for objectivists if formal factors were sufficient to determine an objective evidential support relation. For then, it would be very plausible that the evidential support relation would meet the requirement of *a priority*, for the same reason that deductive logic is *a priori*.

But unfortunately, formal factors aren't enough. Goodman (1955) gave an example which convinced many that the prospects for such an inductive logic were dim. Start with the datum that observing 1,000 green emeralds and no non-green ones is strong

evidence that all emeralds are green. So an inductive logic should assign ‘All emeralds are green’ a high degree of evidential support, relative to the evidence ‘Emeralds 1-1000 are all green, and they are all the emeralds we have observed.’

Now introduce a new predicate ‘grue’ which applies to an object if and only if that object is green and has been observed or else is blue and has not been observed.⁷ Crucially, any observed emerald that is green is also grue. If we then replace all occurrences of ‘green’ with ‘grue’ in the evidence and hypothesis sentences mentioned above, our inductive logic must assign a high degree of evidential support to ‘All emeralds are grue’ relative to the evidence ‘Emeralds 1-1000 are all grue, and they are all the emeralds we have observed.’ But this is plainly the wrong result. Observations of grue emeralds (and lack of observation of any non-grue emeralds) don’t support the claim that all emeralds are grue to anything like the degree to which observations of green emeralds (and lack of observation of any non-green ones) support the claim that all emeralds are green.⁸ In Goodman’s terms, ‘green’ is a projectible predicate whereas ‘grue’ is not.

As Titelbaum (2010, 479) observes, Goodman’s example is suggestive but isn’t an airtight proof that no inductive logic is possible. For instance, it might be that once we take into account the other sentences that will be among the agent’s evidence in any natural spelling-out of Goodman’s case, facts about the syntactic features of these other evidence sentences will suffice to secure the result that ‘green’ is projectible while ‘grue’ is not.

However, in an important result, Titelbaum (2010; see also his 2011) proves that Goodman’s language-dependence problem will infect any purely formal theory of evi-

⁷This definition of ‘grue’ is slightly different from Goodman’s. His ‘grue’ applies to an object just in case it is green and was first observed before t (for some unspecified time t) or blue and first observed after t . I drop reference to t for simplicity.

⁸It would go too far to say that observations of grue emeralds don’t support the claim that all emeralds are grue *at all*. After all, such observations rule out various hypotheses which are incompatible with the hypothesis that all emeralds are grue. For instance, they rule out the hypothesis that all emeralds are red. So they must support the claim that all emeralds are grue to at least some small degree.

dential support. With some plausible auxiliary assumptions, he proves that no evidential support relation can be both substantive (in the sense of sometimes favoring one hypothesis over another despite being logically independent of both) and formal (in the sense of being invariant under permutations of predicates). In order to be substantive, the evidential support relation must be sensitive to more than just the syntactic form of the sentences in question. It must play favorites among predicates.

Some philosophers might immediately jump to the stronger conclusion that there can be no substantive evidential support relation at all. They make the background assumption that such a relation must be purely formal if it's to be worth any salt. Such philosophers might be motivated by two different concerns. The first is epistemic. If the evidential support relation can't be characterized formally, how can we know about it? I won't go into detail about the concern here (after all, much of the remainder of the paper deals with this epistemic issue) except to note that in ethics, most theorists take it that there can't be a formal characterization of moral reasons, but they don't become nihilists or error theorists on that account. Rather, they espouse a wide variety of metaethical positions—naturalism, non-naturalism, expressivism, and the like—that each give subtle answers to the question of how we can know about moral reasons absent any formal characterization of them. The second is a worry about normative force. Facts about evidential support are supposed to constrain what it's rational to believe. But some philosophers are skeptical about the normative force of any rational requirement that can't be motivated by something akin to a Dutch Book argument showing that if you violate the requirement, then there are cases where your attitudes license you to act in a way that is to your own acknowledged, predictable disadvantage. And as it happens, only purely formal rational requirements appear to be amenable to motivation via a Dutch Book argument. Hence if the evidential support relation isn't purely formal, it's doubtful whether it can have any normative force. I think such

philosophers are mistaken, and indeed I am skeptical of whether Dutch Book arguments themselves are prohibitive (see my 2013, 2015a), though again I won't argue for this point here. Rather, I'll just take my target audience to be people who aren't already skeptical of the possibility of a substantive (i.e. non-formal) rational requirement having normative force.

The challenge for objectivists about evidential support is therefore to say what substantive factor or factors determine the evidential support relation, subject to the constraint that they be *a priori*.

A quick caveat, though, before plunging ahead. While I will be exploring approaches on which there is just a single substantive factor that does all the work of determining the evidential support relation, it is also open to the objectivist to be a primitivist or a pluralist about evidential support. On a primitivist view, facts about evidential support are brute facts that aren't determined by other, explanatorily prior facts. Of course, we would still be able to give rules of thumb or *ceteris paribus* principles (e.g. simpler hypotheses tend to be better supported by the evidence) for getting a grip on the evidential support relation, but there aren't further facts that ground the evidential support relation. On a pluralist view, there are multiple substantive factors that together ground the evidential support relation. For instance, facts about simplicity, explanatoriness, beauty, naturalness, and perhaps others besides are given weights and then aggregated to determine the degree to which a given body of total evidence supports some hypothesis. Both primitivism and pluralism strike me as eminently reasonable positions. But before retreating to either of these positions, I want to examine the prospects for a more monolithic⁹ conception of objective evidential support. Still, it is important to remember that these are very live options. Even if a monolithic view of evidential support is implausible in the end, objectivism is not yet defeated.

⁹I owe this term to Jonathan Weisberg, who recently used it in a lecture on objective Bayesianism.

3 Naturalness

Start with perhaps the most popular proposed solution to the Goodman’s grue paradox. This view posits a distinction between natural properties and unnatural ones, and, more generally, posits degrees of naturalness for properties. Natural properties are ones that carve nature at its joints, so that sharing a natural property makes for objective similarity, whereas sharing an unnatural property needn’t do so. In Goodman’s case, greenness is a fairly natural property, while grueness is a horribly unnatural one. So here is a proposal about evidential support more broadly: the substantive factors that determine the evidential support relation are the degrees of naturalness of the properties involved in the hypotheses in question. In general, hypotheses with natural properties received higher degrees of evidential support than hypotheses with unnatural properties.¹⁰

Lewis (1983) posits a distinction between natural and unnatural properties to do a host of jobs, from helping to fix the laws of nature to contributing to determining the referents of our terms. But for present purposes, the most important role of naturalness that Lewis discusses involves its connection to rationality, and in particular its potential

¹⁰Titelbaum (2010) holds that his proof shows that the naturalness proposal, or something very close to it, is a necessary condition for objectivism to be true. The reason is that if a substantive evidential support relation cannot treat evidential support as invariant under permutations of predicates, it can seem that any such evidential support relation must treat some predicates (or the properties to which they refer) as special, whether we call this specialness ‘naturalness’ or something else. But in my view, the fact that the evidential support relation cannot be invariant under permutations of predicates does not entail that it must treat some predicates (or properties) as special *in and of themselves*, such that the appearance of certain predicates in the evidence and hypothesis always makes for greater evidential support. Instead, it could be that the evidential support relation cares about the *relationship* between the predicates and the object terms of which they are predicated; that is, it must care about the meaning of the proposition as a whole (and about the relationship between the evidence proposition and the hypothesis proposition). This is akin to the Moorean (1903) notion of an organic unity. Consider the fact that chocolate and strawberries (together) taste better than chocolate and salmon. This doesn’t mean that the presence of strawberries *in and of itself* makes for a better dish than does the presence of salmon. After all, salmon and greens tastes better than strawberries and greens. It’s the relationship between the two that matters. For this reason, naturalness isn’t the only game in town when it comes to objective evidential support. In fact, as I will argue shortly, it isn’t even the most promising approach.

to solve Goodman's grue paradox. Lewis holds that part of the reason we in fact have beliefs about greenness rather than grueness (at least until we've read our Goodman) is that beliefs that involve projecting greenness are more rational than beliefs that involve projecting grueness, and this in turn is because greenness is more natural than grueness. Here are a couple representative passages:

The principles of charity will impute a bias toward believing that things are green rather than grue, toward having a basic desire for long life rather than for long-life-unless-one-was-born-on-Monday-and-in-that-case-life-for-an-even-number-of-weeks. (Lewis 1983, 375)

We think that some sorts of belief and desire... would be unreasonable in a strong sense... utterly unintelligible and nonsensical. Think of the man who, for no special reason, expects unexamined emeralds to be grue... What makes the perversely twisted assignment of content incorrect, however well it fits the subjects behaviour, is exactly that it assigns ineligible, unreasonable content when a more eligible assignment would have fit behaviour equally well. (Lewis 1986, 38-9)

Lewis' solution to Goodman's paradox, and the broader theory of rationality it yields, has generated widespread interest, much of it sympathetic, in the literature. Here is Jenkins (2013, 100) discussing the use of naturalness in responding to Goodman's paradox, though I hasten to add that she does not ultimately endorse the proposal:

Empirical evidence alone could be construed as underdetermining whether all emeralds are green or all emeralds are grue, but [a theory of rationality on which beliefs involving projecting natural properties are more rational] holds out hope of a solution, provided that it can be argued that the proposition *s*: *All emeralds are green* is more natural than *r*: *All emeralds are grue*.

Similarly, Weatherson argues that we should view naturalness as playing a role in determining reference (so that our words tend to refer to natural rather than unnatural properties) in virtue of the role it plays in the theory of rationality and hence in the

determination of mental content. His is an exegetical essay on Lewis' own theory, but Weatherson also endorses the interpretation of Lewis that he favors:

But why is it more charitable to attribute beliefs about greenness to beliefs about grueness? I think it is because we need more evidence to rationally form a belief that some class of things are all grue than we need to form a belief that everything in that class is green. And that's because, *ceteris paribus*, we need more evidence to rationally form a belief that all *F*s are *G*s than that all *F*s are *H*s when *G* is less natural than *H*. (Weatherson 2012, 3)

Goodman and Titelbaum have argued that an evidential support relation has to play favorites among properties, but if we believe in a distinction between natural and unnatural properties, this should come as no great surprise. The evidential support relation should play favorites: it should favor natural properties over unnatural ones.

3.1 Is Naturalness *A Priori*?

While the naturalness theorist seems to have a nice solution to the problem of language dependence, Titelbaum (2010) argues that it ultimately fails. This is because he thinks that it's an empirical matter what the natural properties are (and, more generally, which properties are more natural than which). If that's right, then his epistemic circularity worry looms: Beliefs about what the natural properties are must be determined on the basis of our evidence, but they aren't themselves entailed by our evidence. If facts about what one's evidence supports are determined in part by facts about what the natural properties are, this means that in order to figure out what one's evidence supports, one must first figure out what the natural properties are, and to do that, one must figure out what one's evidence supports, but to do that one must figure out what the natural properties are. And round and round we go.¹¹

¹¹As Titelbaum (2010, 485) puts the point, the problem is that any substantive evidential support relation 'displays a bias towards certain properties that is prior to and independent of the influence

Now, we might resist this circularity worry by holding that facts about naturalness can determine what you ought to believe even if you're not in a position to know these facts.¹² But even so, facts about naturalness will fail to supervene on one's mental states, at least assuming that the empirical status of facts about naturalness means they're also contingent. The particular supervenience worry dissolves if they are necessary *a posteriori*, though I think it would still threaten the plausibility of the claim that one has default justification for employing naturalness-based evidential support relation in arriving at one's beliefs.

How about the second reason for thinking that facts about evidential support must be *a priori*, namely that inductive methods (or, in our terminology, evidential support relations) must be immodest? Well, if the proposal is that (to use Goodman's phrase) one ought to project the properties that are in fact natural, regardless of whether your

of any evidence. But supplying such a bias was the job the natural properties were supposed to do! In order for the list of natural properties to play its envisioned role in shaping the evidence favoring relation, it cannot be determinable from an agent's evidence.'

One might attempt to invoke reference magnetism to escape Titelbaum's circle. According to one interpretation of the idea of reference magnetism, it is a brute constraint on interpreting the meanings of terms in a language that they tend to refer to natural rather than unnatural properties. So, if you just project the predicates that are already in your language, you're likely to be projecting a natural property. Now, Titelbaum worries that there is still a problem, since despite a general preference for natural properties to serve as the referents of predicates, one's language is still likely to contain plenty of predicates referring to unnatural properties as well. So one will still need empirical evidence to know which of the predicates in one's language to project. Of course, the naturalness theorist might respond that in any case *most* of one's predicates will refer to natural ones due to the general constraints imposed by reference magnetism, and so by adopting a general policy of projecting any and all predicates in one's language, one makes it likely that in any particular instance one will be projecting a natural properties.

Even so, there is a deeper problem with this escape route. For arguably, it is in virtue of facts about rationality, and in particular about evidential support, that predicates in one's language tend to refer to natural rather than unnatural properties. Weatherson (2012) makes a compelling case that this is the order of explanation in Lewis' philosophy of language, at least. There, the idea is that facts about mental content are explanatorily prior to facts about linguistic meaning, and facts about mental contents are themselves determined in part by facts about rationality. It's easier to have rational attitudes toward propositions involving natural properties, so there's pressure to interpret agents as having attitudes toward such propositions rather than propositions involving unnatural properties. And that's why linguistic predicates tend to refer to natural rather than unnatural properties. But this Lewisian picture is one we can't endorse if we hold that the fact that one's terms tend to refer to natural properties is itself what explains how and why it is rational to project natural properties.

¹²Titelbaum considers this possibility in Section 5 of his 2010.

evidence supports the hypothesis that those are the natural properties, then the naturalness theorist's evidential support relation needn't be immodest. But if the proposal allows for defeaters, so that while you ought to start out by projecting the properties that are in fact natural, you should not longer project them if you later get misleading evidence that they aren't natural, then it won't be immodest.

In any event, it is possible to resist the thought that facts about naturalness are empirical. To preview, my suggestion will be that it is *a priori* which properties, if instantiated, would be natural (or which possible properties are more natural than which), even though it's clearly empirical which properties are actually instantiated. To use Hawthorne and Dorr's (2013) helpful way of putting it, the view is that it is *a priori* which properties are possibly natural, and being possibly natural and actually instantiated is necessary and sufficient for being natural. (More generally, it is *a priori* whether some property is possibly natural-to-such-and-such-degree, and being possibly natural-to-such-and-such-degree and actually instantiated is necessary and sufficient for being natural-to-such-and-such-degree.) However, I will ultimately argue (Section 3.2) that the question of the empirical status of naturalness is a red herring, for naturalness is not the most promising substantive factor for determining the evidential support relation. Therefore, the impatient reader can skip ahead without losing the thread.

Start with why one might think it's empirical what the natural properties are (ignoring for the moment whether the empiricity lies in whether a properties would be natural if instantiated, or just in whether a property is in fact instantiated). Certainly, prominent proponents of the natural/unnatural distinction emphasize the empirical nature of naturalness. Lewis and Armstrong both take the sciences, and physics in particular, to be guides to what the natural properties are. For instance, Armstrong (1979, 8), who takes natural properties to be the *only* properties there are, writes, 'What properties and relations there are in the world is to be decided by total science,

that is, by the sum total of all enquiries into the nature of things.’ And he emphasizes the special role of physics: ‘...I look for the genuine universals, the genuine properties and relations, in the fundamental notions of physics’ (1988, 105-6).¹³ In the same vein, Lewis (1984, 228) holds that ‘physics discovers which things and classes are the most elite [i.e. natural] of all’ and takes the most plausible version of the naturalness thesis to be ‘one that give[s] a special elite status to the ‘fundamental physical properties’: mass, charge, quark colour and flavour.’¹⁴ And van Fraassen (1989), a prominent critic of the natural/unnatural distinction, likewise assumes that empirical evidence can be our only guide to what the supposedly natural properties are: ‘if laws are to be what science hopes to provide in the end, then science had better hope to formulate its theories in a correct language [i.e. one whose predicates refer to natural properties]. And the guardians of this correctness can only be the scientists themselves.’

This empirical conception of naturalness has stuck. In a recent critical survey, Hawthorne and Dorr (2013, 18) identify empiricism as one of the central roles played by the notion of naturalness: ‘The right method for identifying actually-instantiated perfectly natural properties is empirical.’ Even authors who question the privileged role assigned to fundamental physics hold that it is an empirical matter what the natural properties are. For instance, Schaffer (2004) holds that all sciences are in the business of discovering the natural properties.

So much for appeals to authority. What is the actual motivation for thinking that it is an empirical matter what the natural properties are? Apart from a background Quinean holism, the main motivation seems to me to be that a central theoretical role of natural properties is to constitute a minimal supervenience base, and it is plausible that it is an empirical matter which properties constitute such a minimal supervenience

¹³These two quotes are highlighted in Schaffer (2004).

¹⁴Another representative quote: Lewis (1986, 60) writes that ‘What physics has undertaken...is an inventory of the *sparse* [i.e. natural] properties of this-worldly things.’

base.

But it is not clear that having natural properties constitute a minimal supervenience base requires it to be an empirical matter whether some property is possibly natural, as opposed to just an empirical matter whether some given possibly natural property is in fact instantiated. Natural properties could constitute a minimal supervenience base in each world even if it was a different set of natural properties in each world, provided just that all the possibly natural properties in each world where they are instantiated are part of a minimal supervenience base for that world.¹⁵

Other roles to be played by natural properties seem to outright conflict with thinking of science, and especially physics, as discovering what the natural properties are. I will mention just two. First is independence: perfectly natural properties are supposed to be mutually independent, in the sense that it is metaphysically possible for any combination of them to be instantiated by some object. Lewis (2009, 209) espouses a combinatorial principle on which every recombination of parts of reality is possible. The relevant parts of reality ‘include not only spatiotemporal parts, but also abstract parts—specifically, the fundamental [i.e. natural] properties.’ The problem is that the properties figuring in fundamental physics don’t seem to be mutually independent. For instance, in the above quote, Lewis mentions mass as a plausible perfectly natural property in virtue of its role in physics. Presumably, the property in question is not (or not only) *having mass*, but rather (or rather also) the infinitely many determinates of that determinable: *having mass 0.298743 grams*, and the like. For part of the naturalness role is supposed to be that perfect duplicates share all there perfectly natural properties, and perfect duplicates cannot have different masses (Hawthorne and Dorr 2013, 44). But these determinate mass properties are not mutually independent. Sim-

¹⁵There may also be grounds for ditching or demoting the supervenience role of the natural/unnatural distinction. For as Schaffer (2004, 99) argues, there may be possible worlds in which there is no minimal supervenience base, and instead, ‘properties might be endlessly supervenient upon lower-level properties.’

ilarly, if there turn out to be multiple fundamental particles, as seems like, and if each fundamental particle corresponds to a perfectly natural property (*being a quark, being a lepton, being an anti-quark, being an anti-lepton*, and the like), then the perfectly natural properties won't be mutually independent.

Second, consider the role supposed to be played by naturalness in determining laws of nature. One might think that this clearly requires that naturalness be empirical, but not so fast. Here is Lewis's take on naturalness and laws. He endorses a Mill/Ramsey-style best system account of laws of nature. A system is a set of propositions (or alternatively, sentences) purporting to describe the world. A suitable system is:

one that has the virtues we aspire to in our own theory-building, and that has them to the greatest extent possible given the way the world is. It must be entirely true; it must be closed under strict implication; it must be as simple in axiomatisation as it can be without sacrificing too much information content; and it must have as much information content as it can have without sacrificing too much simplicity. A law is any regularity that earns inclusion in the ideal system. (Lewis 1983, 367)

However, Lewis notes that an immediate problem arises:

Different ways to express the same content, using different vocabulary, will differ in simplicity. The problem can be put in two ways, depending on whether we take our systems as consisting of propositions (classes of worlds) or as consisting of interpreted sentences. In the first case, the problem is that a single system has different degrees of simplicity relative to different linguistic formulations. In the second case, the problem is that equivalent systems, strictly implying the very same regularities, may differ in their simplicity. (ibid)

Lewis's solution is to bring in naturalness:

We should ask how candidate systems compare in simplicity when each is formulated in the simplest eligible way; or, if we count different formulations as different systems, we should dismiss the ineligible ones from candidacy.

An appropriate standard of eligibility [is] not far to seek: let the primitive vocabulary that appears in the axioms refer only to perfectly natural properties...[Then] laws will tend to be regularities involving natural properties. Fundamental laws, those that the ideal system takes as axiomatic, must concern perfectly natural properties. (ibid)

I take it that science, and especially physics, is supposed to discover the (perfectly) natural properties by discovering the laws of nature, and then we take the natural properties to be those that figure in those laws. But then a potential circularity looms. What the natural properties are depends on what the laws of nature are (if they are whatever properties figure in those laws), but what the laws of nature are depends on what the natural properties are (since laws are regularities in the ideal system, and ideality depends in part on simplicity when the system is formulated in a language whose predicates refer to natural properties). If naturalness is supposed to help settle which system is ideal, and hence what the laws are, then the naturalness facts must be prior to facts about what the laws are and hence cannot hold in virtue of facts about what the laws are. Of course, naturalness could play this role in helping determine what the laws are if it's *a priori* which properties are possibly natural but *a posteriori* which possibly natural properties are actually instantiated.

So the theoretical roles meant to be played by naturalness do not seem to require, and in some cases conflict with, the view that it's an empirical matter which properties are possibly natural, and not just empirical which possibly natural properties are instantiated. Thinking that it's an empirical, or at least a contingent, matter which properties are possibly natural also raises some uncomfortable questions:

Among the properties which are not in fact perfectly natural, which ones could be perfectly natural? One possible view is that every property whatsoever could have been perfectly natural. While it seems like an attraction of this view that it enables us to dodge the need for making a distinction between the possibly perfectly natural properties and the rest, plenty of other awkward questions remain. For example, assuming that it is still necessary

that no perfectly natural property supervenes on all the rest, there must be limits on which sets of properties can be perfectly natural *together*; and it is hard to think of a good way to answer questions like ‘Which properties are such that they could be perfectly natural in a world where *being a spoon* was perfectly natural?’ (Hawthorne and Dorr 2013, 32-33)

As I noted above, there is room in logical space for the view that it is necessary but *a posteriori* which properties are possibly natural (to-such-and-such-degree). But I think the most attractive view on which such facts are necessary is one on which they come out *a priori* as well.¹⁶ First, note that for non-perfectly-natural properties, it is already plausible that it is *a priori* which are more natural than which. For instance, it is plausible that it is *a priori* that greenness is more natural than grueness. So the tough cases concern scientific theoretical properties. But here, it is tempting to adopt a view on which, were the laws of nature much different—or perhaps even *any* different—from how they are now, the actually instantiated properties figuring in our best scientific theories wouldn’t be instantiated at all. If massive bodies didn’t attract (ignoring other forces), then they wouldn’t count as massive (and more generally, no mass properties would be instantiated). If protons weren’t composed of quarks, they wouldn’t be protons (and the property of being a proton wouldn’t be instantiated). And if quarks were composed of smaller particles, then they wouldn’t be quarks (and the property of being a quark wouldn’t be instantiated).

On this view, it is plausibly *a priori* that if the property *being a proton* is instantiated at all, then it’s natural (more generally, natural-to-such-and-such-degree), since in order for that property to be instantiated, the scientific theory in which it figures must be true, and this is all that’s required in order for it to count as natural.

¹⁶Lewis himself certainly thought that it was a non-contingent matter which properties are possibly natural, though it is unclear whether he thought this was also *a priori*. He writes that, ‘The name [natural] has proved to have a drawback: it suggests to some people that it is supposed to be *nature* that distinguishes the natural properties from the rest; and therefore that the distinction is a contingent matter, so that a property might be natural at one world but not at another. I do not mean to suggest any such thing. A property is natural or unnatural *simpliciter*, not relative to one or another world’ (1986, 60, fn 44).

If this is right, then we could have our objective evidential support relation involve the constraint that one only project possibly natural properties (which may or may not be actually instantiated).¹⁷ No circularity worry arises, since if it is *a priori* whether a property is possibly natural, then one needn't do induction in order to find out what properties are possibly natural, and hence in order to find out how to do induction.

3.2 The Shortcomings of Naturalness

However, in the final analysis I think that the question of whether naturalness is an empirical matter is a bit of a red herring. This is because naturalness is narrowly targeted at enumerative induction, but cannot be the whole story in determining the privileged evidential standards more broadly. While it is reasonably clear to see how an appeal to natural properties would help solve Goodman's grue paradox, and perhaps even the more general problem of justifying *enumerative* induction, it is doubtful whether it would be of much use in justifying induction more broadly.

Why can't naturalness be the whole story? Jenkins (2013, 102) points out that naturalness seems impotent to solve the problem of skepticism about the external world. In order for it to do so, it would have to be the case that the property of being such that one's perceptions are veridical counts as a more natural property than the property of being a brain in a vat being fed non-veridical perceptual experiences (or something along those lines). But insofar as I have a grip on the distinction between natural and unnatural properties in the first place, I don't find myself with any strong judgment either way about which property is more natural than the other. So insofar as one

¹⁷Of course, if one cannot refer to uninstantiated but possibly natural properties, this complicates matters. Take one of Titelbaum's central examples, that of phlogiston theory at a time when it was still a live possibility. On some views, phlogiston necessarily doesn't exist. It wouldn't exist even if the Ramsey sentence corresponding to phlogiston theory were true. If that is right, the same difficulties arise in modeling the doxastic states of phlogiston theories and interpreting their language that arise for fictional terms like 'Santa Claus.' I won't go into these matter here.

thinks of the problem of skepticism about the external world as a kind of skepticism about induction, namely about inductive reasoning from evidence concerning perceptual experiences to hypotheses about the external world, naturalness can't be the whole story about induction.

Now, this might not be so bad, for there are attractive solutions to the problem of skepticism about the external world whose diagnosis is that beliefs about the external world aren't inductive after all. To give just one example, consider Williamson's (2000) thesis that your evidence consists of all and only the propositions that you know ($E=K$). Assuming that you know that you have hands (or some other external world proposition), then by believing that you have hands you are not going beyond what's entailed by your evidence. On this diagnosis, Cartesian skepticism about the external world is a different problem, and admits of a different solution, than Humean skepticism about induction.

But even setting aside Cartesian skepticism, appeals to naturalness fall short of what is required to ground an objective evidential support relation. Consider cases of non-enumerative induction (sometimes called 'abduction'). Seeing that the sky is filled with dark clouds, I conclude that it will rain shortly. How is naturalness supposed to be relevant here? The property of being such that it will rain isn't obviously any more natural than the property of being such that it won't rain. (Similarly for the other properties you might think are involved in this case—the property of being such that there are dark clouds and it will rain isn't clearly more natural than the property of being such that there are dark clouds and it won't rain.) Then there are the sorts of inductive inferences common in the sciences. From the evidence that the organism have such-and-such phenotypic property, the scientist concludes that it probably is heterozygous at a given locus rather than homozygous for the dominant allele. But the property of being heterozygous is not more natural than being homozygous for the

dominant allele.

If we're worried about induction and evidential support in general, rather than just with enumerative induction and Goodman's grue paradox, then an appeal to naturalness can't be the full story, even if it can avoid Titelbaum's circularity objection.¹⁸ Of course, this leaves open the possibility that an appeal to naturalness might be part of the story of what grounds facts about evidential support. But in the next section I will argue that there is a more promising approach to the problem of underdetermination that, if successful, would render appeals to naturalness superfluous.

4 Explanationism

My proposed alternative solution is a familiar one: explanationism. Facts about what evidence supports what hypotheses to what degree ultimately turn on explanatory considerations. Explanationism tends to go under the name 'inference to the best explanation,' but this name might misleadingly suggest that evidence can only support a hypothesis when the hypothesis is the best possible explanation of the evidence. But a given body of evidence can also support a hypothesis in cases where the hypothesis, if true, would be well explained by the evidence. This would be a sort of 'inference from the best explanation,' to use Lipton's (1991) term. More generally, two propositions can each be evidence for the other, even though explanation is asymmetric. Nevertheless, these observations are compatible with the claim that there is an objective evidential support relation, and that this relation is determined by explanatory considerations.

¹⁸A caveat: Hawthorne and Dorr (2013) mention the possibility of evaluating entire probability functions for naturalness, with the result that facts about evidential support boil down to facts about which probability functions are most natural. I won't pursue this possibility here.

4.1 Explanation as a General Response

Explanationism can succeed where appeals to naturalness failed. It has the potential to account not only for how not only intuitively good instances of enumerative induction can be rational, but also for how good instances of induction more generally can be rational. My discussion here will be programmatic. I won't be saying nearly enough to demonstrate that explanationist treatments of all cases of induction are correct. Instead, I just want to show that explanationism is promising as a general solution to the problem of induction, rather than being narrowly aimed at enumerative induction.

First, explanationism may be able to serve as a solution to the problem of skepticism about the external world. Bonjour has argued that the best explanation of the coherence and stability of our perceptual beliefs is that the experiences they are based on are by and large veridical:

The coherence-cum-stability of a system of beliefs is complicated and fragile, easily disrupted or destroyed, and thus it is inherently unlikely that a system of beliefs which is constantly receiving [observational inputs] would remain coherent from moment to moment without constant revisions which would destroy its stability. Some explanation is therefore needed for why it continues to do so, and the obvious one is that the beliefs of the system match the independent reality which they purport to describe closely enough to minimize the potential for disruptive input. (Bonjour 1985, 171)

I don't want to hang my hat on Bonjour's proposal (see Vogel 2010 for criticism of Bonjour's argument). Note also that even if the general explanationist thrust of his proposal is right, we might dispute the details. For instance, Bonjour takes the relevant explanandum to be the long-run coherence and stability of our perceptual beliefs. An alternative explanationist solution to the skeptical challenge might take the explananda to be propositions like *that I am having an experience as of a hand in front of me* and then to argue that the best explanation for why I'm having this particular experience is that there is in fact a hand in front of me.

Even if an explanationist solution to skepticism about the external world fails, this needn't be terribly worrying, for as I have noted, there are other solutions on offer, such as Williamson's E=K thesis, that treat beliefs about the external world as not inductive at all (at least in the good case).

More importantly, explanatory considerations are clearly relevant to garden variety cases of non-enumerative induction where appeals to naturalness are impotent. In the case of the cloudy sky, the best explanation for why there are dark clouds covering the sky would also be a good explanation of why it will rain (if in fact it does rain). Here we are inferring to the best explanation of the dark clouds, and then inferring from that best explanation to the conclusion that it will rain. Similarly, insofar as our scientist's inference is a good one, this is plausibly because being heterozygous at the relevant locus is a better explanation for having the given phenotype than is being homozygous for the dominant allele.

Explanatory considerations also have the potential to provide an attractive solution to Goodman's grue paradox, which would render appeals to naturalness superfluous. White (2005b) proposes that one explanatory virtue is stability: an explanation of some fact F is stable to the extent that, according to this explanation, F couldn't easily have failed to obtain, and stability is a good-making feature of an explanation to the extent that its explanandum calls out for explanation.

You've observed a bunch of emeralds, and all of them have been green. By definition, this means that all observed emeralds have also been grue. Nonetheless, White argues that the hypothesis that all emeralds are green is a better explanation of the fact that all observed emeralds have been green/grue than is the hypothesis that all emeralds are grue. The reason is that the all-green hypothesis is more stable than the all-grue hypothesis. Properties like grueness are 'counterfactually dependent on observation,' whereas ordinary color properties are not:

Whereas an unobserved green thing would still have been green had we observed it, an unobserved *grue* thing would have been *bleen* [i.e. blue if observed and green otherwise] had we observed it. For if something is *grue* and unobserved then it is blue, and would have been observed and blue and hence *bleen*. Similarly, an unobserved *bleen* thing would have been *grue* had we observed it. For if something is *bleen* and unobserved then it is green and would have been green had we observed it, in which case it would have been green and observed and hence *grue*. (White 2005b, 18)¹⁹

Importantly, it was arbitrary which emeralds we happen to have observed.²⁰ Had we happened to have observed different emeralds, then some actually unobserved emeralds would have been observed. The hypothesis that all emeralds are green has the consequence that even had we observed some of the actually unobserved emeralds, our explanandum—that all observed emeralds are green/*grue*—would still have obtained.

¹⁹Note that the fact that gruesome properties are counterfactually dependent on observation is independent of whether we start with non-gruesome predicates and define gruesome ones therefrom, or if we instead start with gruesome predicates and define non-gruesome ones therefrom. Even if we define ‘green’ as ‘*grue* if observed and *bleen* if unobserved,’ it remains the case that an unobserved green thing (which is therefore *bleen*) would still have been green had we observed it (in which case it would have been *grue*).

²⁰What about cases where it isn’t wholly arbitrary which emeralds we’ve observed? Consider a different property *grue+*, where something is *grue+* just in case it is green and first observed before AD 2500 or blue and first observed (if ever) after AD 2500. Suppose also that for any given emerald we don’t observe AD 2500, the reason it doesn’t get observed until then is that it’s on an exoplanet, and we don’t get the technology to mine on exoplanets until after AD 2500. Then, the hypothesis that all emeralds are *grue+* provides a fairly stable explanation of why we’ve observed only *grue+* emeralds. If all emeralds are *grue+*, then even if we’d made somewhat different observations than we in fact did, this wouldn’t involve our mining on exoplanets in the 21st century, and so we still would have observed only *grue+* emeralds. I think there are three things to say about this sort of case.

First, even if the hypothesis that all emeralds are *grue+* is a rather stable explanation of our evidence, it is still a bit a less stable explanation than the hypothesis all emeralds are green. If the former hypothesis is true, then it is still possible for us not to have gotten the evidence that we’ve observed only *grue+* emeralds, even if we have to go to fairly dissimilar possible worlds to get cases where we lack this evidence. So the *grue+* hypothesis may be a lot better than the standard *grue* hypothesis in terms of explanatory stability, but it’s still not as good as the green hypothesis. Second, we know that by around the late 25th century, we (i.e. later humans) will start heavily favoring the green hypothesis over the *grue+* hypothesis (and White’s story explains why this would be justified), so reflection-style reasoning suggests we should favor the green hypothesis now. Third, even if there are gruesome properties where White’s explanationist story breaks down (consider the property *grue++*, which applies to things that are green unless they’re in some causally inaccessible region of the universe), we can still fall back on an appeal to naturalness, provided that it’s *a priori* which properties are possibly natural, as I have suggested. This would amount to a retreat from a monolithic picture of evidential support, on which explanatory considerations do all the work, to a pluralist approach on which there are a variety of substantive factors that together determine the evidential support relation.

By contrast, the hypothesis that all emeralds are grue doesn't have this consequence. If all emeralds are grue, then had we observed some actually unobserved emeralds, then it wouldn't be true that all observed emeralds are green/grue, for these emeralds would be blue/bleen.

There is of course much more to be said here, and I refer the reader to White's paper for further discussion. But I do think that White's explanationist response to the grue paradox is on the right track, and if that's right, then appeals to naturalness are superfluous, and explanatory considerations have the potential to be the whole story in grounding facts about evidential support. This would be a monolithic picture of evidential support which yields a treatment of enumerative induction as just a special case of induction more generally (though see footnote 20 for a caveat on which naturalness may still have a role to play). In this I am simply echoing the line taken in Harman's (1965, 88) seminal paper on inference to the best explanation: 'all warranted inferences which may be described as instances of enumerative induction must also be described as instances of the inference to the best explanation.'

4.2 Explanationism and the A Priori

So far we have seen that explanationism has the potential to be a response to worries about underdetermination in general, whereas appeals to naturalness seem to help only with enumerative induction. Now I want to argue that an explanationist approach to evidential support avoids the circularity objection that Titelbaum levels against the naturalness approach. For it is plausible that the most basic explanatory standards are *a priori*. (I don't say '*our* most basic explanatory standards,' for it is doubtful whether the standards that we fallible humans have adopted are the correct ones.)

Importantly, this is not to say that whether something is in fact a good explanation of a given body of evidence is an *a priori* matter. For starters, something can't count as

a good explanation of anything unless it is true (or at least approximately true), and this will almost always be an empirical matter (mathematical explanations being perhaps the lone example, which we can set aside since we're concerned with induction). Moreover, on many familiar accounts of explanation, it will be an empirical matter whether a given set of propositions, even if true, would constitute an explanation of the evidence in question (Ludlow 1991, 60). For instance, on the deductive-nomological model (Hempel and Oppenheim 1948), an explanation must include a lawlike statement, but even holding fixed that the statement is true, it will be a contingent, empirical matter whether it is a law or instead an accident. Similarly, on causal accounts of explanation, even granting that some event occurred (so that the corresponding proposition about its occurrence is true), it will be a contingent, empirical matter whether it caused the evidence. Better to say, following Ludlow, that a potential explanation of some evidence is one that does explain that evidence in some possible world or other.

One might also think that it is an empirical matter what the best potential explanation of some body of evidence will be, since which potential explanation looks best will depend on our background knowledge, much of which is empirical. For instance, it is an empirical matter whether the best potential explanation of the sidewalk's being wet is that it recently rained or that someone washed it down with a hose. But this is compatible with thinking that it is an *a priori* matter which hypothesis is the best potential explanation of a body of *total* evidence.

My suggestion, then, is that it is an *a priori* matter what the privileged explanatory standards are, and hence what the best potential explanation of a given body of total evidence is. We have default justification for employing those explanatory standards, and those explanatory standards are not defeasible on empirical grounds.

I won't attempt to give an analysis of explanatory goodness. Indeed, I doubt whether such an analysis is possible. But as a first pass, the most basic explanatory virtues,

the features that make a potential explanation a good one, include the extent to which the potential explanation is simple,²¹ unified (if this is different from simplicity), and stable (in White’s sense, above).

In suggesting that the most basic explanatory standards are *a priori*, I want to make two points. First, unlike facts about what the natural properties are, there is little pull toward thinking of physics (or the sciences in general) as being in the business of discovering facts about the importance of simplicity, unification, stability, and the like in judging the goodness of a potential explanation. Science discovers how simple the world in fact is, but not how important simplicity is to a good explanation. Indeed, if we adopt an expressivist approach to epistemology,²² there won’t be any such facts at all. (A caveat: expressivists are typically keen to earn the right to talk just like realists and hence to be able to talk of facts about the relevant normative domain, but these ‘facts’ won’t be the sort of facts that are discoverable by scientific means.)

Second, it is difficult to come up with a case in which the most basic explanatory standards should be revised on empirical grounds. Where there is a case of methodological change on empirical grounds, it seems unlikely that the jettisoned principle was a *basic* explanatory principle in the first place. This is to be expected, of course, if explanationism is the whole story about evidential support, since any adequate evidential support relation must be immodest and hence treat itself as unrevisable on empirical grounds.

Take an example from the history of science. Prior to Newton, scientific standards

²¹One might worry that facts about simplicity will depend on facts about naturalness, for the reason given by Lewis in his discussion of laws of nature. If simplicity is a matter of, say, the number of symbols needed to state the hypothesis, then how simple a hypothesis is will depend on the language used to express it. So we need to fix on a canonical language in which to evaluate the simplicity of different hypotheses, and one option would be for the canonical language to be one in which the predicates refer to natural properties. But if facts about which properties are possibly natural-to-such-and-such-degree are *a priori*, as suggested above, then this doesn’t threaten to make facts about simplicity, and hence explanatoriness, empirical. An alternative option, which I won’t pursue here, would be to conceive of simplicity in non-syntactic terms.

²²See Field (2009) for a defense.

had it that appeals to action at a distance were a cost of a scientific theory, even a major cost. But afterwards, the methodological principle of avoiding appeals to action at a distance was dropped. Now, a first thing to note is that it's not clear that this methodological principle was dropped on *empirical* grounds, as opposed to conceptual ones. Perhaps it was Newton's conceptual progress in formulating an elegant, mathematically precise theory involving action that predicted the extant evidence at least as well as its competitors that led to the dropping of this methodological principle, rather than any new empirical evidence. But regardless of whether the methodological principle was dropped on empirical or conceptual grounds, the preference against theories involving action at a distance is not plausibly regarded as having enjoyed the status of a basic principle of inductive methodology. Better to view the case of Newton and action at a distance as one where more basic explanatory standards led to the dropping of a non-basic explanatory principle. The fact that Newton's theory was simple and provided a unified treatment of terrestrial and heavenly motion (favoring simplicity and unification being plausible components of our basic explanatory standards), and that this theory involved action at a distance, meant that the non-basic explanatory principle proscribing appeals to action at a distance had to be dropped.

Here is a purely hypothetical, example from Field that makes the same point:

Presumably our empirical methodology includes a bias for simplicity. We recognize that insofar as we can account for all past and present observations by our present body of theory T, we could account for it equally well by an alternative T* according to which T holds until Jan 1 [2016], after which Aristotelian physics, Lamarckian biology, etc., take over. Why do we rule out T*, and base our predictions instead on the approximate truth of T? ...presumably it's that T is a vastly simpler way of accommodating our evidence than is T*. But now it might seem that our methodology of choosing the simpler is empirically revisable (either by revising the principle 'choose the simpler' or by revising the simplicity judgments that give this slogan its content). Suppose we had evidence that in each past year on New Year's day, the laws of nature drastically changed; that would seem like good

inductive evidence that they'd change on New Years in [2016] too. Doesn't this show that our empirical methodology...is itself empirically revisable?

No, it doesn't show this at all. What it shows is only that we regard theories T^{**} according to which the laws of nature change every year as more plausible than corresponding theories T^{***} according to which the laws change every year until 2004, but don't change then. It seems that we have two pre-existing biases: one for T over T^* , which licenses a belief that the laws won't change in [2016] given evidence that they haven't changed in the past; the other for T^{**} over T^{***} , which would license a belief that the laws will change in [2016] were we to be given evidence that they have changed in the past. So the fact that the laws of nature haven't changed drastically in the past is indeed inductive evidence that they won't change drastically in [2016]; but this fact is based on a fixed bias (for T over T^* and for T^{**} over T^{***}) which there is no obvious way to undermine by empirical evidence. (Field 2005, 75)

And even if we could find a case where we really do revise the weight we assign to simplicity (or unification, or stability) on empirical grounds, this might just be taken to show that simplicity wasn't a basic explanatory or inductive virtue. Some more basic explanatory principles stand in the background and govern how the weights assigned to simplicity, unification, and stability should change given changes in evidence.²³

So here's why a broadly explanationist treatment of evidential support doesn't fall prey to a version of Titelbaum's circularity objection. First, the most basic explanatory standards—favoring simplicity, unification, stability, and the like—aren't the sorts of things that science is in the business of discovering. Science employs such explanatory standards, but doesn't discover facts about which are the right ones. Second, it's

²³Perhaps no evidence could simultaneously call into question the preference for simplicity, unification, and stability (or whatever one's candidate basic explanatory standards are). In that case, at least one such explanatory standard would be treated as empirically indefeasible in that particular situation, and would be used to determine how to modify the methodological principles that are called into question. This would amount to a picture on which at least one methodological principle is treated as *a priori* relative to a given body of evidence, even though no methodological principle is treated as *a priori* relative to all possible bodies of evidence. (Compare Bueno and Colyvan (2004) on the law of contradiction in logic.) Of course, provided that there is a uniquely rational way to determine which principles to treat as indefeasible for present purposes, and how to revise those principles that are called into question, it is tempting to say that really, the most basic methodological principles are those that determine this uniquely rational way to proceed with methodological change, even if we find it difficult or impossible to actually articulate those most basic principles.

not clear how to come up with a case in which the most basic explanatory standards should be revised on empirical grounds. This is for good reason—a good inductive method must be immodest and hence treat its most basic methodological principles as empirically indefeasible. If the most basic methodological principles are explanatory ones, this would explain why the most basic explanatory standards, whatever they are, aren't subject to revision on empirical grounds.

5 Conclusion

The evidential support relation cannot be characterized in purely formal terms. This is the lesson of Goodman's grue paradox and Titelbaum's related proof. So if there is such a thing as evidential support, there must be substantive (i.e. non-formal) constraints on what counts as evidence for what. But these substantive constraints must be *a priori* ones in order for evidential support to play its role in determining what one ought to believe. Some philosophers have proposed that these substantive constraints are to be found in the distinction between natural and unnatural properties: only natural properties are projectible. But even we can defend the claim that it's *a priori* which properties are possibly natural, naturalness can't be the whole story about evidential support. Better to adopt an explanationist approach, which has the potential to handle not just enumerative induction, but evidential support more broadly. Such an explanationist approach avoids Titelbaum's circularity objection, since the most basic facts about explanatoriness are *a priori*.

If successful, this defense of objective explanatory support is an essential part of defending the claim that all requirements of epistemic rationality are synchronic. For objective explanatory support is necessary for defending Uniqueness, and Uniqueness is plausibly necessary for defending such a purely synchronic picture of epistemic ratio-

nality, on which the only norm is one stating that your beliefs at a time ought to be the ones that are supported by your evidence at that time.

References

- Armstrong, David. 1979. *A Theory of Universals*. Cambridge: Cambridge University Press.
- Bueno, Otávio, and Mark Colyvan. 2004. 'Logical Non-Apriorism and the Law of Non-Contradiction.' In Priest, Beall, and Armour-Garb (eds). *The Law of Non-Contradiction: New Philosophical Essays*. Oxford University Press.
- Carnap, Rudolf. 1950. *The Logical Foundations of Probability*. Chicago: University of Chicago Press.
- Colyvan, Mark, Regan, Helen, and Scott Ferson. 2001. 'Is it a Crime to Belong to a Reference Class?' *The Journal of Political Philosophy* 9 (2): 168–81.
- Conee, Earl, and Richard Feldman. 2001. 'Internalism Defended.' In *Epistemology: Internalism and Externalism*, edited by Hilary Kornblith, 231-260. Oxford: Blackwell Press.
- Elga, Adam. 2010. 'How to Disagree about How to Disagree.' In *Disagreement*, edited by Richard Feldman and Ted Warfield, 175-186. New York: Oxford University Press.
- Field, Hartry. 2000. 'Apriority as an Evaluative Notion.' In *New Essays on the A Priori*, edited by Paul Boghossian and Christopher Peacocke, 117-149. New York: Oxford University Press.
- Field, Hartry. 2005. 'Recent Debates about the A Priori.' In *Oxford Studies in Epistemology*, Vol 1, edited by Tamar Szabo Gendler and John Hawthorne, 69-88 Oxford: Oxford University Press.
- Field, Hartry. 2009. 'Epistemology without Metaphysics.' *Philosophical Studies* 143 (2): 249–90.
- Goodman, Nelson. 1955. *Fact, Fiction, and Forecast*. Cambridge, MA: Harvard University Press.
- Harman, Gilbert. 1965. 'The Inference to the Best Explanation.' *The Philosophical Review* 74 (1): 88–95.
- Harman, Gilbert. 1986. *Change in View*. Cambridge, MA: MIT Press.
- Hawthorne, John, and Cian Dorr. 2013. 'Naturalness.' In *Oxford Studies in Metaphysics*, Vol 8, edited by Karen Bennett and Dean Zimmerman, 3-77. Oxford: Oxford University Press.
- Hedden, Brian. 2013. 'Incoherence without Exploitability.' *Noûs* 47 (3): 482–95.
- Hedden, Brian. 2015a. 'Options and Diachronic Tragedy.' *Philosophy and Phenomeno-*

logical Research 90: 423–51.

Hedden, Brian. 2015b. ‘Time-Slice Rationality.’ *Mind* 124 (494): 449–91.

Hedden, Brian. 2015c. *Reasons without Persons: Rationality, Identity, and Time*. Oxford: Oxford University Press.

Hempel, Carl and Paul Oppenheim. 1948. ‘Studies in the Logic of Explanation.’ *Philosophy of Science* 15 (2): 135–175.

Hlobil, Ulf. 2015. ‘There are Diachronic Norms of Rationality.’ *Thought* 4 (1): 38–45.

Horowitz, Sophie. 2013. ‘Immoderately Rational.’ *Philosophical Studies* 167 (1): 1–16.

Jenkins, Carrie. 2013. ‘Justification Magnets.’ *Philosophical Studies* 164 (1): 93–111.

Lewis, David. 1971. ‘Immodest Inductive Methods.’ *Philosophy of Science* 38 (1): 54–63.

Lewis, David. 1983. ‘New Work for a Theory of Universals.’ *Australasian Journal of Philosophy* 61(4):343–377.

Lewis, David. 2009. ‘Ramseyan Humility.’ In *Conceptual Analysis and Philosophical Naturalism*, edited by David Braddon-Mitchell and Robert Nola, 203–222. Cambridge, MA: MIT Press.

Lipton, Peter. 1991. *Inference to the Best Explanation*. New York: Routledge.

Meacham, Christopher. 2014. ‘Impermissive Bayesianism.’ *Erkenntnis* 79 (6): 1185–1217.

Moore, G.E. 1903. *Principia Ethica*. Cambridge: Cambridge University Press.

Moss, Sarah. 2014. ‘Credal Dilemmas.’ *Noûs* (early view), doi 10.1111/nous.12073

Podgorski, Abelard. ‘A Reply to the Synchronist.’ Forthcoming in *Mind*.

Schaffer, Jonathan. 2004. ‘Two Conceptions of Sparse Properties.’ *Pacific Philosophical Quarterly* 85(1): 92–102.

Titelbaum, Michael. 2010. ‘Not Enough There There: Evidence, Reasons, and Language Independence.’ *Philosophical Perspectives* 24 (1): 477–528.

Titelbaum, Michael. 2011. ‘Symmetry and Evidential Support.’ *Symmetry* 3: 680–98.

Weisberg, Jonathan. 2009. ‘Locating IBE in the Bayesian Framework.’ *Synthese* 167 (1): 125–143.

White, Roger. 2005a. ‘Epistemic Permissiveness.’ *Philosophical Perspectives* 19 (1): 445–459.

White, Roger. 2005b. ‘Explanation as a Guide to Induction.’ *Philosophers’ Imprint* 5

(2): 1–29.