

Can All-Accuracy Accounts Justify Evidential Norms?

Christopher J. G. Meacham

Forthcoming in *Epistemic Consequentialism*.

Abstract

Some of the most interesting recent work in formal epistemology has focused on developing accuracy-based approaches to justifying Bayesian norms. These approaches are interesting not only because they offer new ways to justify these norms, but because they potentially offer a way to justify all of these norms by appeal to a single, attractive epistemic goal: having accurate beliefs. Recently, Easwaran & Fitelson (2012) have raised worries regarding whether such “all-accuracy” or “purely alethic” approaches can accommodate and justify evidential Bayesian norms. In response, proponents of purely alethic approaches, such as Pettigrew (2013b) and Joyce (2016), have argued that scoring rule arguments provide us with compatible and purely alethic justifications for the traditional Bayesian norms, including evidential norms. In this paper I raise several challenges to this claim. First, I argue that many of the justifications these scoring rule arguments provide are not compatible. Second, I raise worries for the claim that these scoring rule arguments provide purely alethic justifications. Third, I turn to assess the more general question of whether purely alethic justifications for evidential norms are even possible, and argue that, without making some contentious assumptions, they are not. Fourth, I raise some further worries for the possibility of providing purely alethic justifications for content-sensitive evidential norms, like the Principal Principle.

1 Introduction

Some of the most interesting recent work in formal epistemology has focused on the project of developing scoring rule arguments for Bayesian norms.¹ These scoring rule arguments apply the apparatus of decision theory to the epistemic realm, appealing to epistemic analogs of decision theoretic principles in order to show that we should satisfy these norms.

¹The term “Bayesian” is often applied to views that endorse both Probabilism (the claim that a subject’s credences should be probabilistic) and Conditionalization (the claim that a subject’s new credences, upon receiving evidence, should be equal to her prior credences conditional on her evidence). Here I use the term “Bayesian norm” broadly, to include any norm that imposes a constraint on a subject’s credences or degrees of belief. Thus Probabilism and Conditionalization are examples of what I’ll call “Bayesian norms”, as are norms like Lewis’s (1980) Principal Principle, van Fraassen’s (1995) Reflection Principle, Indifference Principles, and the like.

One reason this project is interesting is that it provides a justification for Bayesian norms which avoids many of the problems facing older attempts, such as dutch book arguments and representation theorem arguments. A second reason this project is interesting is that it has been seen as providing a way of justifying Bayesian norms that's compatible with *Veritism* — the view that the only epistemic goal is to have beliefs that are as accurate as possible, and every epistemic norm must ultimately boil down to and be justified in terms of this goal.² The thought is that if we can plug in some measure of accuracy into our decision theoretic principles, and then show how it follows from this that we should satisfy various Bayesian norms, then we'll be able to provide a *purely alethic justification* for Bayesian norms — a justification that doesn't appeal to anything besides the goal of having accurate beliefs. And if we can provide purely alethic justifications for all of the main Bayesian norms, then we have a way of justifying the Bayesian framework that is congruous with Veritism.

Reconciling Bayesianism and Veritism would be a remarkable achievement. For a number of the central Bayesian norms are *evidential norms*, norms whose prescriptions depend on what one's evidence is. And, as Easwaran & Fitelson (2012) note, Veritism seems to be in tension with the existence of evidential norms. According to Veritism, the only thing we should care about (epistemically speaking) is having accurate beliefs. But the existence of evidential norms suggests that we should also care about having beliefs that line up with the evidence. And it seems like these two desiderata can come apart: a subject can have beliefs that line up with the evidence but are nonetheless false. So if accuracy is the only thing we should care about, then it looks like we should give up on norms which care about other things as well, like evidential norms.

In response to these kinds of worries, Pettigrew (2013b) and Joyce (2016) have argued that scoring rule arguments can provide us with an array of purely alethic justifications for the main Bayesian norms, including evidential norms. And they've maintained, at least implicitly, that these justifications are compatible with one another. Indeed, Pettigrew (2013b) has argued that not only can scoring rule arguments provide us with such justifications, existing scoring rule arguments *do* provide us with such justifications. I'll call this stance on the justificatory status of the Bayesian norms the "All-Accuracy Account":

The All-Accuracy Account: There are scoring rule arguments that provide us with compatible purely alethic justifications for the main Bayesian norms, including evidential norms like the Principal Principle, the Indifference Principle, and Conditionalization.

In this paper I'll raise several worries for the All-Accuracy Account. I'll raise some worries about whether the existing scoring rule arguments provide justifications for these norms that are compatible and purely alethic. And I'll raise some concerns regarding whether it's even possible to provide purely alethic justifications for evidential Bayesian norms.

Despite the critical nature of this paper, it should not be seen as an attempt to bury scoring rule arguments. Like many others, I think that the investigation of scoring rule arguments is one of the most important recent developments in formal epistemology, and that the work done by Joyce, Pettigrew and others in exploring this framework is invaluable. Rather, this paper

²For a defense of this view, see Goldman (2002).

should be understood as an attempt to get clear on what these arguments actually establish, one which suggests an appropriately modest stance regarding these accomplishments. Or, if one prefers, this paper can be understood as an attempt to provide a helpful “to do” list of challenges for proponents of views like the All-Accuracy Account to address.

The rest of this paper will proceed as follows. In the second section I’ll provide some background on scoring rule arguments, and the dominance principles that these arguments employ to derive various Bayesian norms. In the third section I’ll assess the claim that the scoring rule arguments that have been offered for these different Bayesian norms are compatible. I’ll argue that many of the standard scoring rule arguments rely on implausible dominance principles, and that when these principles are fixed up, many of these arguments are no longer compatible. In the fourth section I’ll assess the claim that the scoring rule arguments that have been offered for these different Bayesian norms provide purely alethic justifications. I’ll argue that the most explicit argument that’s been offered for this conclusion is unsound. In the fifth section I’ll consider whether purely alethic justifications for evidential norms are possible. I’ll suggest that such justifications are not possible unless one adopts contentious versions of these norms which many would reject. In the sixth section I’ll suggest that even if one adopts a contentious version of a norm like the Principal Principle, it’s not possible to provide a purely alethic justification for such a norm. In the seventh section I’ll conclude by sketching some big picture morals.

2 Background

In this section, I’ll lay out some background regarding scoring rule arguments. There are a number of ways to develop such arguments. My presentation here will largely follow Pettigrew’s, as he’s presented perhaps the most detailed attempt to provide a unified set of scoring rule arguments for the Bayesian norms.³ In what follows I’ll lay out some preliminary terminology, briefly discuss the epistemic utility functions these arguments employ, and sketch the decision theoretic principles these arguments employ in order to obtain the various Bayesian norms.

2.1 Preliminaries

Let F be the *doxastic propositions* of a subject; that is, the set of propositions about which a subject has beliefs. Let W be the set of consistent truth assignments to these propositions. I’ll call the elements of W the *worlds* of a subject.⁴ I’ll assume that F is closed under the usual logical operations (conjunction, disjunction, negation, etc). Following Pettigrew (2013b), I’ll assume that F and W are finite.

³See Leitgeb & Pettigrew (2010a), Leitgeb & Pettigrew (2010b), Pettigrew (2012), Pettigrew (2013a), Pettigrew (2013b), and Pettigrew (2014).

⁴These “worlds” will not, of course, be as fine-grained as possible worlds, which correspond to consistent assignments of truth values to *every* proposition (regardless of whether anyone has beliefs about them or not).

Let a subject’s *credence function* be a function cr which takes propositions as arguments, and yields real numbers in the $[0,1]$ interval, with higher numbers representing greater degrees of confidence. Let C be the set of all of the possible credence functions the subject might assign.

2.2 Epistemic Utility

Let an *epistemic utility function* be a function U which takes a credence function and a world as arguments, and yields a real number representing the epistemic utility of holding that credence function at that world.

Proponents of Veritism, such as Pettigrew (2013b) and Joyce (2016), maintain that satisfactory epistemic utility functions must be “alethic”, and must assign quantities that value accuracy and nothing else. As such, they require satisfactory epistemic utility functions to satisfy certain constraints. For example, since they want accuracy to be an overriding factor, they require epistemic utility functions to satisfy:⁵

Truth-Directedness: If for all propositions A , $cr(A) \geq cr^*(A)$ if A is true at w and $cr(A) \leq cr^*(A)$ if A is false at w , and for some A , A is true at w and $cr(A) > cr^*(A)$, or A is false at w and $cr(A) < cr^*(A)$, then $U(cr, w) > U(cr^*, w)$.

Likewise, since they want accuracy to be the *only* relevant factor – and thus want factors such as the content of propositions to be irrelevant — they require epistemic utility functions to satisfy:⁶

Extensionality: Let $f_{w,w^*} : F \rightarrow F$ be a bijection between true and false propositions at w with true and false propositions at w^* ; i.e. if $f_{w,w^*}(A) = B$, then either A is true at w and B is true at w^* , or A is false at w and B is false at w^* . If cr and cr^* are such that, for all A , $cr(A) = cr^*(f_{w,w^*}(A))$, then $U(cr, w) = U(cr^*, w^*)$.

These are the two most prominent constraints on alethic utility functions. Some proponents of Veritism have suggested that alethic utility functions should satisfy some other constraints as well, though there is disagreement regarding what exactly these further constraints should be.

Fortunately, we needn’t hash out the details of this debate here. For there is one epistemic utility function which is widely agreed to be adequate, namely, the epistemic utility function associated with the Brier score.⁷ Let the *indicator function* be a function I which takes a world w and a proposition A as arguments and yields either 0 (if A is false at w) or 1 (if A is true at w). Then the epistemic utility function associated with the Brier score is:⁸

$$U(cr, w) = - \sum_{A \in F} (I(w, A) - cr(A))^2. \quad (1)$$

⁵See Joyce (2009).

⁶See Joyce (2009).

⁷Though see Levinstein (2012) for a reason to be wary of the Brier score.

⁸The Brier score itself is what you get if you remove the minus sign: it’s a measure of *inaccuracy*, such that the higher the values it yields, the less accurate the credence function.

For the purposes of this paper, I'll follow Pettigrew (2013b) in assuming that this is the correct epistemic utility function.⁹

2.3 Dominance Norms and Bayesian Norms

Scoring rule arguments have been offered for a number of different Bayesian norms.¹⁰ I'll largely focus here on the arguments offered by Pettigrew. These arguments generally appeal to a dominance norm of some sort, and proceed by using this dominance norm and an epistemic utility function to derive the relevant Bayesian norm.

Let's see how these arguments go. In what follows, let F be the set of doxastic propositions of the subject; W the worlds associated with F ; cr , cr^* and cr^{**} credence functions defined over F ; and U the correct epistemic utility function.

One of the core Bayesian norms is *Probabilism*:

Probabilism: A subject's credences should be probabilistic.

The scoring rule arguments for Probabilism employ something like *Strict Dominance* as a dominance norm:

Definition: A function cr is *strictly dominated* by cr^* for a subject *iff*, for all $w \in W$, $U(cr, w) < U(cr^*, w)$.

Strict Dominance: If a subject is such that (i) cr is strictly dominated by cr^* , and (ii) cr^* is not strictly dominated by any cr^{**} , then it is impermissible for that subject to adopt cr .

There's some variation in the literature regarding the exact form of the dominance norm employed in the argument for Probabilism. Some employ a slightly weaker notion of dominance, some omit the second clause, and some employ one notion of dominance in the first clause and a slightly different notion of dominance in the second.¹¹ But for our purposes, these differences won't matter.

Another core Bayesian norm is *Conditionalization*:^{12,13}

⁹Some, like Joyce (2016), hold that there is no unique correct epistemic utility function. Instead, there's a small group of "kosher" utility functions, any one of which is good enough. But proponents of such views still hold that the utility function associated with the Brier score is one of the functions that's good enough; and that's sufficient for our purposes.

¹⁰For example, see Joyce (1998), Greaves & Wallace (2006), Joyce (2009), Leitgeb & Pettigrew (2010a), Leitgeb & Pettigrew (2010b), Pettigrew (2012), Easwaran (2013), Pettigrew (2013a), Pettigrew (2013b), and Pettigrew (2014).

¹¹For example, see Joyce (2009) for instances of the first and second variations, and Pettigrew (2013b) for an instance of the third. I present Strict Dominance in the text (as opposed to one of these variants) in order to keep the form of the different dominance norms I present in this section as similar as possible.

¹²The formulation of Conditionalization given here is a standard one. Pettigrew (2013b) applies the label "Conditionalization" to a different kind of norm — a synchronic norm regarding what plans one should make, not a diachronic norm about what credences one should have. One might motivate moving from a diachronic norm to a synchronic norm in various ways (e.g., see Meacham (2010), Hedden (2015), and Moss (forthcoming)). But since the synchronic rule described in Pettigrew (2013b) diverges quite a bit from the diachronic rule Bayesians generally endorse (and the rule discussed by in other places, such as Leitgeb & Pettigrew (2010a)), I've stuck with the standard formulation here.

¹³Like most formulations of Conditionalization, this formulation leaves open a number of questions regarding how

Conditionalization: If a subject with probabilistic credences cr receives evidence E , her new credences cr_E should be:

$$cr_E(A) = cr(A | E), \text{ if defined.} \quad (2)$$

The scoring rule arguments for Conditionalization need to employ something like what I'll call *Credal Expected E-Dominance*:¹⁴

Definition: Let the subject's credences be cr_0 . Assume cr_0 is probabilistic, and let cr and cr^* be probability functions. Then cr is *credal expected E-dominated* by cr^* for that subject iff:

$$\sum_{w \in E} cr_0(w) \cdot U(cr, w) < \sum_{w \in E} cr_0(w) \cdot U(cr^*, w). \quad (3)$$

Credal Expected E-Dominance: If a subject is such that the next piece of evidence she'll receive is E at t' , and is such that (i) cr is credal expected E -dominated by cr^* , and (ii) cr^* is not credal expected E -dominated by any cr^{**} , then it is impermissible for that subject to adopt cr at t' .

Some formulations of the decision-theoretic norms this argument requires obscure the diachronic nature of these norms, and the norm's appeal to the subject's evidence.¹⁵ But these features of the norm are crucial to getting the derivation to go through. As one might expect, you can't get diachronic results about what credences a subject should have given certain evidence without building diachronic features and appeals to evidence into one's decision theoretic norm.

In addition to Probabilism and Conditionalization, most Bayesians accept a chance-credence principle of some kind. There is some disagreement about how to best formulate this principle, but one popular formulation, borrowing from Lewis (1980), is:¹⁶

Principal Principle: Let ch_G be the chance function picked out by G . The credences of a subject with total evidence G , cr_G , should be such that:

$$cr_G(A) = ch_G(A), \text{ if defined.} \quad (4)$$

exactly the norm should be understood (see Schoenfield (2016) and Meacham (2016) for discussions of some of these issues). For example, should we understand cr as the subject's credences after she's received all of her evidence except E , the most recent piece of evidence she's received, or should we understand cr as a credence function at some arbitrary time, and E as the cumulative evidence the subject's received since then? But the ambiguities of this formulation are largely orthogonal to the issues I'll be raising in this paper, so I'll put these complications aside.

¹⁴This is, roughly, a dominance-style formulation of the decision theoretic norm employed by Leitgeb & Pettigrew (2010a) in their scoring rule argument for Conditionalization. But there are a number of other ways to formulate and divide up the norms the argument requires; see, for example, Easwaran (2013). (There are also a number of other interesting issues regarding these arguments that I lack the space to discuss; see Schoenfield (2016).)

¹⁵E.g., see Pettigrew (2013b) and the formulation of Maximize Subjective Expected Utility.

¹⁶This formulation follows from the conjunction of Lewis's formulation and Conditionalization. Lewis's initial formulation only makes appeal to a subject's initial credences, prior to the receipt of evidence.

Pettigrew’s (2013a) scoring rule argument for the Principal Principle employs what I’ll call *Chancy Expected Dominance*:

Definition: Let C be the set of possible chance functions. Let E be the total evidence of the subject. A function cr is *chancy expected dominated* by cr^* for that subject *iff*, for all $ch \in C$:

$$\sum_{w \in W} ch(w | E) \cdot U(cr, w) < \sum_{w \in W} ch(w | E) \cdot U(cr^*, w). \quad (5)$$

Chancy Expected Dominance: If a subject is such that (i) cr is chancy expected dominated by cr^* , and (ii) cr^* is not chancy expected dominated by any cr^{**} , then it is impermissible for that subject to adopt cr .

It’s worth noting that Chancy Expected Dominance’s appeal to the subject’s evidence is crucial to the success of the argument. After all, we want the permissible credences regarding chance events to *vary* depending on what a subject’s evidence is. (E.g., if a subject has no evidence other than that the chance of heads is 1/2, their credence in heads should be 1/2. But if the subject has no evidence other than that (i) the chance of heads is 1/2 *and* (ii) the coin landed heads, their credence in heads should be 1.) And one can’t derive a norm whose prescriptions are evidence-dependent, like the Principal Principle, without employing a decision theoretic norm that appeals to evidence.¹⁷

Finally, some Bayesians endorse a fourth norm, an Indifference Principle of some kind. There are a number of different kinds of Indifference Principles, but one popular version of this principle, considered by Pettigrew (2014), is:¹⁸

Indifference Principle: The credences of a subject with total evidence E , cr_E , should be such that:

$$\forall w_i, w_j \in E, cr_E(w_i) = cr_E(w_j). \quad (6)$$

Pettigrew’s (2014) scoring rule argument for this Indifference Principle employs *Worst-Case Dominance*:

Definition: Let E be the total evidence of a subject. A function cr is *worst-case dominated* by cr^* for that subject *iff* there is some $w \in E$ such that, for all $w^* \in E$, $U(cr, w) < U(cr^*, w^*)$.

Worst-Case Dominance: If a subject is such that (i) cr is worst-case dominated by cr^* , and (ii) cr^* is not worst-case dominated by any cr^{**} , then it is impermissible for that subject to adopt cr .

¹⁷At first glance the version of the dominance norm used to derive the Principal Principle given in Pettigrew (2013b) doesn’t seem to make any claims about evidence. But it needs to be understood as implicitly making claims of this kind (as the version presented in Pettigrew (2013a) does) in order for the argument to go through.

¹⁸This is the principle that Pettigrew (2014) calls “PoI+”, but the same general points apply with respect to what Pettigrew calls “PoI”. In both cases, the principle’s prescriptions are sensitive to what the subject’s evidence is (in the latter case, requiring the subject to not have any evidence), and in both cases the derivation of the principle won’t follow unless we employ a norm in which the subject’s evidence plays a role.

Again, some presentations of the scoring rule argument for the Indifference Principle employ formulations of these decision-theoretic norms that obscure their appeal to the subject's total evidence.¹⁹ But these appeals to the subject's total evidence are a vital part of these norms. For without employing such a norm, the derivation of the Indifference Principle won't go through.

3 Compatibility

Proponents of the All-Accuracy Account maintain that scoring rule arguments provide us with compatible and purely alethic justifications for the main Bayesian norms. One worry for this claim, which we'll focus on in this section, has to do with whether the scoring rule arguments for these different norms are compatible.

Pettigrew is the most prolific producer of scoring rule arguments for Bayesian norms, as well as one of the most careful, so I'll focus on his arguments here. The scoring rule arguments for Bayesian norms generally take the following form: first they argue that we should adopt a certain kind of epistemic utility function, then they argue that we should adopt a certain decision-theoretic norm, and then they show how to use this utility function and decision-theoretic norm to derive the relevant Bayesian norm. As we saw in section 2.3, the decision-theoretic principles Pettigrew employs are generally dominance principles of the following form:²⁰

X-Dominance: If a subject is such that (i) cr is X -dominated by cr^* , and (ii) cr^* is not X -dominated by any cr^{**} , then it is impermissible for that subject to adopt cr .

By plugging in different notions of dominance into this schema, we get different decision-theoretic principles. And we can use these different principles to derive the different Bayesian norms.

I'll argue that this method of deriving Bayesian norms is problematic. First, I'll suggest that these kinds of dominance principles are only plausible in isolation. When considered jointly, these principles become implausible. Second, I'll note that even though we can get around this worry by modifying the dominance principles these derivations employ, doing so requires us to reassess these derivations. In particular, I'll show that once we've chosen which forms of dominance we want to be epistemically relevant, we can formulate principles tracking these notions of dominance that *are* jointly plausible. But once we adopt these modified principles, we need to reassess these derivations of Bayesian norms. For in many cases these derivations will no longer be compatible.

To get a feel for how this problem arises, let's start by considering why Pettigrew formulates these dominance principles using the above schema. It's clear why the schema includes

¹⁹E.g., the version of the dominance norm used to derive the Indifference Principle given in Pettigrew (2013b) doesn't seem to make any claims about evidence. But it needs to be understood as implicitly making claims of this kind (as the version endorsed in Pettigrew (2014) does — see sections 4 and 6 for the key caveats and revisions) in order for the argument to go through.

²⁰Though the form is a bit different in the case of diachronic dominance principles (such as Credal Expected E-Dominance).

the first clause — the thought that there’s something *prima facie* wrong with dominated options is a key part of the intuition underwriting these dominance norms. But why does this schema include the second clause? At first glance, anyway, it’s tempting to simply drop the second clause and adopt a schema like this:

X-Dominance (v1): If a subject is such that cr is X -dominated by cr^* , then it is impermissible for that subject to adopt cr .

Pettigrew’s (2013b) reason for rejecting the Dominance (v1) schema is that it leads to absurd results in cases in which there are no undominated credences. If every credence function is dominated by another credence function, then Dominance (v1) will entail that every credence function is impermissible. If we take epistemic dilemmas — situations in which every epistemic option is impermissible — to be impossible, then this is a *reductio* of Dominance (v1).²¹

One way to avoid this problem is to modify the Dominance (v1) schema so that being dominated only rules out a credence function if there are undominated credence functions to be had. This yields the following schema:

X-Dominance (v2): If a subject is such that (i) cr is X -dominated by cr^* , and (ii) there are undominated credence functions, then it is impermissible for that subject to adopt cr .

But, as Pettigrew (2013b) notes, this revised schema is still problematic. For both Dominance (v1) and (v2) yield the wrong results in certain cases. Here is a slight variant of a case Pettigrew (2013b) considers. Suppose a subject can adopt any one of a countably infinite sequence of credence functions defined over two worlds, w_1 and w_2 . And suppose the epistemic utilities of adopting each of these credence functions at each world are the values given by the following table:

	cr_0	cr_1	cr_2	...	cr_n	...
w_1	2	1	2	...	n	...
w_2	2	$2 - \frac{1}{2}$	$2 - \frac{1}{4}$...	$2 - \frac{1}{2^n}$...

Given a dominance relation that tracks these epistemic utilities, cr_0 will be the only undominated credence function. The function cr_0 will not be dominated, since it yields a higher

²¹This is a bit quick, for one can avoid a *reductio* by rejecting any of the premises that led to the contradiction, and there are other premises one could reject other than Dominance (v1). In particular, one could reject the premise that there are any normatively relevant notions of dominance such that there will be no undominated options in some cases. But this is an unattractive option: since these schemas are supposed to spell out the general relationship between normatively relevant notions of dominance and what we ought to believe, one would like to avoid having to make substantive assumptions about what dominance relations can hold.

In a similar fashion, one can respond to the counterexample to Dominance (v1) and Dominance (v2) given below by maintaining that there aren’t any normatively relevant notions of dominance that will yield to the pattern of dominance relations the case describes. But again this is an unattractive option, given that these schemas are supposed to spell out the general relationship between normatively relevant notions of dominance and what we ought to believe.

epistemic utility at w_2 than any other credence function. But for every $i > 0$, cr_i will be dominated by cr_{i+1} . And since i is unbounded, it follows that every $cr_{i>0}$ will be dominated.

Given Dominance (v1) or (v2), it follows that cr_0 will be the only permissible credence function. But this doesn't seem right. For example, every probability function that assigns a non-zero value to w_1 will assign some $cr_{i>0}$ a higher expected utility than cr_0 . Thus it seems like at least some $cr_{i>0}$'s should be permissible.

The moral, Pettigrew notes, is that being dominated isn't in itself a bad thing. Whether being dominated is a bad thing depends on the status of what's doing the dominating. If an option is dominated by something which has a mark against it, then this doesn't suffice to rule out the dominated option. Being dominated only rules out an option if it's dominated by something which is itself unimpeachable. Thus, in the example above, the fact that every $cr_{i>0}$ is dominated doesn't rule them out, as every credence function that dominates them is itself dominated by something else.

Call an option *unimpeachable* iff no normatively relevant factor tells against it. If we restrict our attention to cases in which the only normatively relevant factor is X -dominance, then an option will be unimpeachable iff it's not X -dominated. Given Pettigrew's moral — that being dominated only tells against an option if the dominating option is unimpeachable — this suggests the following schema:

X -Dominance (v3): If a subject is such that (i) cr is X -dominated by cr^* , and (ii) cr^* is not X -dominated by any cr^{**} , then it is impermissible for that subject to adopt cr .

This is, of course, precisely the schema Pettigrew employs.

Dominance (v3) works well if we assume that X -dominance is the only normatively relevant factor. But suppose that's not the case, and that there are two normatively relevant factors. Perhaps, in addition to X -dominance, there is also a "side constraint" that directly states that certain options are impermissible.²² Or perhaps there are two normatively relevant notions of dominance, X -dominance and Y -dominance. How should we proceed, in order to determine what options are permissible?

One approach is to apply each of these constraints sequentially. So, for example, one might first apply a side constraint to narrow down the set of permissible options, and then appeal to X -Dominance (v3) to narrow down the set of permissible options some more. Or one might appeal to X -Dominance (v3) to narrow down the set of permissible options, and then appeal to Y -Dominance (v3) to further narrow down the set of permissible options. The problem with this approach, as Easwaran & Fitelson (2012) show, is that the final result can be order-dependent.²³ That is, the set of permissible options one is left with after applying both constraints can depend on the order in which we apply these constraints. This leaves us

²²I'll take a side constraint to be a constraint which directly rules out certain options (as opposed to normatively relevant notions of dominance, which only rule out options indirectly, such as by slotting into a Dominance norm of some kind). Of course, the distinction between normatively relevant factors that are "side constraints" and those that are not is a vague one. I borrow this terminology from Pettigrew (2013b), who himself borrows it from Nozick (1974).

²³Though see Pettigrew's (2014) for an instance of someone willing to bite the bullet here. (In order to resolve the tension between the Principal Principle and the Indifference Principle, Pettigrew suggests that we apply the norms that yield the Principal Principle before we apply those yielding the Indifference Principle.)

in the uncomfortable position of having to decide between different orderings, and of having to provide some rationale for why this is the right ordering to choose.

A second approach is to apply both of these constraints to the entire set of options, and to then take the intersection of the options that remain to be the permissible options. So, for example, one might take the intersection of the options not ruled out by a side constraint, and the options not ruled out by X -Dominance (v3). Or one might take the intersection of the options not ruled out by X -Dominance (v3), and the options not ruled out by Y -Dominance (v3). This approach avoids the ordering problems that afflicted the first approach, since there is no order in which the constraints are applied — they're each applied independently.

But this approach is problematic for the same reasons that Dominance (v1) was problematic. For if, say, X -Dominance (v3) and Y -Dominance (v3) jointly rule out every option, then this approach will entail that every option is impermissible. Assuming that epistemic dilemmas are impossible, this is a *reductio* of this approach.²⁴

A third approach avoids this problem by modifying the schema so that being dominated only rules out a credence function if there are unimpeachable credence functions to be had — credences functions that are undominated with respect to every normatively relevant notion of dominance, compatible with any side-constraints, and so on:

X -Dominance (v4): If a subject is such that (i) cr is X -dominated by cr^* , and (ii) there are unimpeachable credence functions, then it is impermissible for that subject to adopt cr .

Then, as before, one applies constraints to the entire set of options, and takes the intersection of what remains to yield the permissible options.

But both the second approach and third approaches yield the wrong results in some cases. Suppose there are exactly two normatively relevant factors, X -dominance and Y -dominance, where X -dominance tracks X -utility, and Y -dominance tracks Y -utility. And consider three credence functions defined over a pair of worlds w_1 and w_2 , whose utilities at each world are:

²⁴Again, as mentioned in footnote 21, this is a bit quick, for one can avoid this *reductio* in other ways, such as by rejecting the assumption that there are two normatively relevant notions of dominance that will rule out every option in this way. Again, however, this is an unattractive option: since these schemas are supposed to spell out the general relationship between normatively relevant notions of dominance and what we ought to believe, one would like to avoid having to make substantive assumptions about what dominance relations can hold.

In a similar fashion, one can respond to the counterexample given below to both the second and third approaches by maintaining that there are any normatively relevant notions of dominance that will yield to the pattern of dominance relations the case describes. But again, this is an unattractive option given that these schemas are supposed to spell out the general relationship between normatively relevant notions of dominance and what we ought to believe.

		cr_1	cr_2	cr_3
X-utility:	w_1	99	100	101
	w_2	99	100	0
Y-utility:	w_1	100	99	0
	w_2	100	99	101

Evaluating this case using the second approach, we find that cr_3 is not ruled out as it's neither X nor Y -dominated by anything. But cr_1 is impermissible because it's X -dominated by cr_2 , and cr_2 is X -undominated. And cr_2 is impermissible because it's Y -dominated by cr_1 , and cr_1 is Y -undominated.

Evaluating this case using the third approach, we find that cr_3 is again not ruled out because it's neither X nor Y -dominated by anything. (Indeed, since we've stipulated that X and Y -dominance are the only normatively relevant factors, cr_3 is unimpeachable.) But cr_1 is impermissible because it's X -dominated by cr_2 , and there exist unimpeachable credence functions (i.e., cr_3). And cr_2 is impermissible because it's Y -dominated by cr_1 , and there exist unimpeachable credence functions.

So both the second and third approaches entail that cr_3 is the only permissible credence function. But these verdicts don't seem right. For example, *every* probability function will take cr_1 to have either a higher expected X -utility or a higher expected Y -utility than cr_3 . And every probability function which assigns a credence greater than ≈ 0.02 to each world will take cr_1 to have both a higher expected X -utility and a higher expected Y -utility than cr_3 .

In light of this, let's consider a fourth approach. Recall Pettigrew's moral: whether being dominated is a bad thing depends on the status of what's doing the dominating. In particular, being dominated only rules out an option if the dominating option is unimpeachable.

This suggests the following schema:

X -Dominance (v5): If a subject is such that (i) cr is X -dominated by cr^* , and (ii) cr^* is unimpeachable, then it is impermissible for that subject to adopt cr .

And, as before, one applies constraints to the entire set of options, and takes the intersection of what remains to yield the permissible options.²⁵

If we adopt this fourth approach, then in the example above cr_1 , cr_2 and cr_3 will all come out permissible. The function cr_1 won't be ruled out, for while cr_1 is X -dominated by cr_2 , cr_2 isn't unimpeachable, as it's Y -dominated by cr_1 . Likewise, cr_2 won't be ruled out, for while cr_2 is Y -dominated by cr_1 , cr_1 isn't unimpeachable, as it's X -dominated by cr_2 .

²⁵In some places, at least, it seems Pettigrew would endorse this fourth approach (e.g., see his argument in favor of "Dominance-" in section 3 of Pettigrew (2013b), and compare Dominance- to Dominance (v5)). In other places, however, it Pettigrew seems to prefer the first approach (see footnote 23).

I suggest that this fourth approach is the one we should adopt. And Dominance (v5) is the schema we should employ to formulate dominance principles.

The dominance-style derivations of Bayesian norms provided in the literature tend to use Dominance (v1) or (v3)-style principles, not Dominance (v5)-style principles.²⁶ The appeal to Dominance (v3)-style principles is unproblematic if we assume that the notion of dominance being employed is the only normatively relevant factor, as in that case Dominance (v3) and Dominance (v5) are equivalent. But the appeal to Dominance (v3)-style principles *is* problematic if we assume that there are multiple normatively relevant factors — something we're committed to if we want to endorse more than one of these arguments. For if there are multiple normatively relevant factors, it's Dominance (v5), not Dominance (v3), that correctly describes how dominance facts bear on what it's permissible to believe. And once we adopt these Dominance (v5)-style principles, these derivations of the Bayesian norms may no longer go through.

3.1 Two Examples

The dominance principle derivations of Bayesian norms given in the literature tend to use Dominance (v1) and (v3)-style principles. I've argued that we should replace them with Dominance (v5)-style principles. But once we make these replacements, and take all of the normatively relevant factors into account, these derivations of Bayesian norms may no longer go through. To get a feel for how this goes, let's work through two simple examples.

For the first example, suppose there to be exactly two normatively relevant factors, strict dominance and chancy expected dominance. To simplify the discussion, let's assume we're restricting our attention to subjects with no evidence.

Now consider the argument that employs strict dominance to establish Probabilism. It's been shown in the literature that every non-probabilistic credence function is strictly dominated by a credence function which is itself not strictly dominated by anything.²⁷ Given Strict Dominance (v3), this entails that non-probabilistic credence functions are impermissible. But given Strict Dominance (v5), this entailment doesn't follow — to show that every non-probabilistic credence function is impermissible, it needs to be shown that every non-probabilistic credence function is strictly dominated by a credence function which is itself neither strictly dominated *nor* chancy expected dominated. And this is not true. For there are cases in which some non-probabilistic credence functions are only strictly dominated by credence functions which are themselves chancy expected dominated.²⁸ Thus, given these assumptions about what's norma-

²⁶E.g., see Joyce (2009) for versions of Dominance (v1)-style principles, and Pettigrew (2014) and Pettigrew (2013a) for versions of Dominance (v3)-style principles.

²⁷E.g., see Pettigrew (2013c) and the references therein.

²⁸The existence of such cases was shown by Easwaran & Fitelson (2012). Here is a simplified version of such a case, offered by Pettigrew (2013b). Consider a subject with a credence of $3/5$ in both A and $\neg A$, who knows that the chance of A is $3/5$, and knows nothing else about the chances (such as the chance of $\neg A$). As Pettigrew (2013b) shows, the subject's non-probabilistic credences will be strictly dominated by a narrow band of probabilistic credence functions which assign values of around $1/2$ to A and $\neg A$. But since none of these probabilistic credence functions assign $3/5$ to A , all of them are chancy expected dominated.

tively relevant, Strict Dominance (v5) won't entail that all non-probabilistic credence functions are impermissible. And the attempt to establish Probabilism by appealing to strict dominance won't work.

That said, given this particular stance regarding what factors are normatively relevant, the failure of this argument doesn't really matter. For while the argument from strict dominance to Probabilism fails, the argument from chancy expected dominance to the Principal Principle succeeds, and the Principal Principle entails Probabilism. Thus we can replace the argument from strict dominance to Probabilism with an argument from chancy expected dominance to Probabilism.

Let's see how the dominance argument for the Principal Principle goes. Given Chancy Expected Dominance (v5), showing that every Principal Principle-violating (or "PP-violating") credence function is impermissible requires showing that every PP-violating credence function is chancy expected dominated by a credence function which is itself neither chancy expected dominated nor strictly dominated. Pettigrew (2013a) shows that every PP-violating credence function will be chancy expected dominated by a PP-satisfying credence function which is not chancy expected dominated. Furthermore, every PP-satisfying credence function is probabilistic, and no probabilistic credence function is strictly dominated.²⁹ Thus every PP-violating credence function is chancy expected dominated by a credence function which is itself neither chancy expected dominated nor strictly dominated. And so, given these assumptions about what factors are normatively relevant, Chancy Expected Dominance (v5) will entail that all PP-violating credence functions are impermissible. And thus the attempt to establish the Principal Principle by appealing to chancy expected dominance succeeds.

Let's turn to the second example. As before, suppose that there are exactly two normatively relevant factors, but this time let these factors be chancy expected dominance and worst-case dominance. To simplify things, let's again assume that we're restricting our attention to subjects with no evidence.

Let's start by assessing the prospects for a dominance argument for the Indifference Principle. Pettigrew (2014) shows that every Indifference Principle-violating (or "IP-violating") credence function is worst-case dominated by the IP-satisfying credence function which is itself not worst-case dominated. And given Worst Case Dominance (v3), this would entail that violating the Indifference Principle is impermissible. But given Worst Case Dominance (v5), more needs to be shown: it needs to be shown that every IP-violating credence function is worst-case dominated by a credence function which is itself not worst-case dominated *and* not chancy expected dominated. And, given some mild assumptions about chance assignments, this won't be true.³⁰ The only function which is isn't worst-case dominated — the IP-satisfying credence function — will be chancy expected dominated by other credence functions.³¹ Thus, given these assumptions about what's normatively relevant, Worst Case

²⁹E.g., see Joyce (1998).

³⁰For example, the following assumptions suffice: (1) some ur-chance function assigns a different non-zero value to some pair of worlds, and (2) the assignments of ur-chance functions don't overlap; i.e., they assign non-zero values to disjoint sets of worlds. (This second assumption will be disputed by some proponents of Humean accounts of chance. But this argument for the Principal Principle is supposed to be independent of whether one is a Humean.)

³¹One can see this intuitively as follows. Pettigrew's derivation shows that credence functions conditional on E

Dominance (v5) won't entail that violating the Indifference Principle is impermissible. And thus the dominance argument for the Indifference Principle fails.

Likewise, the dominance argument for the Principal Principle fails. To show that violating the Principal Principle is impermissible given Chancy Expected Dominance (v5), it needs to be shown that every PP-violating credence function is chancy expected dominated by a credence function which is itself neither chancy expected dominated nor worst-case dominated. But given the same mild assumptions as before, there won't be any such function — the only functions which aren't chancy expected dominated will be worst-case dominated. So, given these assumptions about what's normatively relevant, Chancy Expected Dominance (v5) won't entail that violating the Principal Principle is impermissible. And thus the argument from chancy expected dominance to the Principal Principle fails. So, given these assumptions about what's normatively relevant, Chancy Expected Dominance (v5) won't entail that violating the Principal Principle is impermissible, and the dominance argument for the Principal Principle will fail.³²

These are two simple examples of how scoring rule arguments for Bayesian norms can break down once we move to Dominance (v5)-style principles, as I've argued we should. In both of these examples there are only two epistemically relevant factors in play. More plausible stances will take there to be more than two epistemically relevant factors.³³ But the general moral is the same for both simple stances and more complicated ones. In order to

that aren't mixtures of the possible chance functions which assign 1 to E will be chancy expected dominated. For subjects with no evidence, it follows that credence functions that aren't mixtures of the possible chance functions will be chancy expected dominated. Now, suppose that (1) some ur-chance function ch assigns a different non-zero value to some pair of worlds, call them w_1 and w_2 , and (2) ur-chance functions don't overlap; i.e., they assign non-zero values to disjoint sets of worlds. Given (1) and (2), the IP-satisfying credence function can't be a mixture of ur-chance functions. For the IP-satisfying credence function will assign the same non-zero value to w_1 and w_2 , while any mixture of ur-chance functions satisfying (1) and (2) must assign w_1 and w_2 either 0 (if ch gets zero weight in the mixture) or distinct non-zero values (if ch gets non-zero weight in the mixture). And since the IP-satisfying credence function can't be a mixture of ur-chance functions, it will be chancy expected dominated.

³²In this case these failures are predictable, since (given the kinds of mild assumptions described in footnote 30) the Principal Principle and Pettigrew's Indifference Principle will yield inconsistent prescriptions. Thus, if one's set up one's framework correctly, it *shouldn't* follow that subjects are obligated to satisfy both the Principal Principle and the Indifference Principle. (The fact that the Dominance (v5) schema yields this result is a mark in favor of it over (say) the Dominance (v1) and (v3) schemas, which will (inconsistently) require subjects to satisfy both.) But even in cases in which these failures are not predictable, the same general points hold. We need to be careful to track all of the normatively relevant factors in play, and see whether the argument goes through given those factors. For the fact that these arguments work in the special case in which there's only one normatively relevant factor is no indication that they'll continue to work once other normatively relevant factors are added into the mix.

³³Of course, there's a sense in which proponents of Veritism will ultimately want the only epistemically relevant factor to be accuracy. But they can still maintain that there are several different kinds of dominance that are all grounded by this concern, and which should all be taken into account when evaluating whether a credence function is permissible. In any case, the main thrust of this section — that the standard scoring rule arguments that have been provided for these Bayesian norms aren't compelling if we take there to be multiple normatively relevant factors in play, and that if we want to endorse more than one of these arguments, we need to show that we can construct analogs of these arguments that take all of the epistemically relevant factors into account — is independent of the question of Veritism.

assess whether an argument for a Bayesian norm goes through, we need to be aware of, and take into consideration, *all* of the normatively relevant factors. And the fact that an argument goes through in the special case in which there's only one normatively relevant factor is no indication that it'll still go through after we introduce other normatively relevant factors.

3.2 Summing Up

Let's take stock. I've argued that the dominance arguments for Bayesian norms should adopt Dominance (v5)-style principles. If this is right, then the dominance arguments for Bayesian norms which appeal to Dominance (v1) or (v3)-style principles are no good. And while arguments employing Dominance (v3)-style principles suffice in cases in which the notion of dominance being employed is the only relevant factor (since in these cases the Dominance (v3) and (v5) schemas are equivalent), these arguments aren't compelling when there are multiple normatively relevant factors in play, as we're required to think if we want to endorse more than one of these arguments at the same time.

Once we realize that we need to adopt Dominance (v5)-style principles, our assessments of the arguments for Bayesian norms becomes a more holistic process. For what a Dominance (v5)-style principle takes to be impermissible will depend on what credence functions are unimpeachable, and thus on what exactly the normatively relevant considerations are. And this makes establishing that (say) it's impermissible to violate the Principal Principle a much trickier task. For this isn't something one can establish on its own, once and for all. Rather, one can only show that, *given that such and such are all of the normatively relevant factors*, it's impermissible to violate the Principal Principle. And any time one wants to take a different stance with respect to what the normatively relevant factors are, the argument needs to be revisited.³⁴

4 Purely Alethic Justification

Advocates of the All-Accuracy Approach take existing scoring rule arguments to provide compatible purely alethic justifications for many Bayesian norms, including evidential Bayesian norms. In the last section we considered whether these arguments provide compatible justifications. In this section, we'll turn to consider whether these arguments provide purely alethic justifications; i.e., justifications that don't appeal to any epistemic value other than that of having accurate beliefs.

³⁴Although this discussion has followed Pettigrew in focusing on dominance-style decision-theoretic principles, I take the general moral to apply regardless of what kinds of decision-theoretic principles are in play. Namely: if there are multiple normatively relevant factors to take into consideration, then a plausible account of what's permissible will be holistic. One can model this holism by appealing to a set of rules that cross-reference each other (as in the text), or by adopting a single decision-theoretic principle which takes all of these factors into account. But either way, one will need to settle on what the normatively relevant factors are before one can demonstrate whether a given option is permissible or impermissible.

Why think that the scoring rule arguments that have been given for these norms yield purely alethic justifications? The most explicit argument for this claim comes from Pettigrew (2013b). Pettigrew replies to Easwaran & Fitelson (2012) worries regarding the compatibility of Joyce's Veritism and evidential norms as follows:

“Joyce’s argument assumes that the only goal of a credal state is accuracy... That is, Joyce’s argument assumes a version of credal virtue monism. ... I wish to defend the monism required for Joyce’s argument... The major stumbling block for this version of credal value monism is that there seem to be evidential norms; and their force seems to stem from the evidential goals at which they demand that we aim. ...we agree that the evidential norms hold, but we show that they follow from the goal of accuracy alone, not from a distinct evidential goal. ...

My strategy is this: Recall the form of Joyce’s argument for Probabilism as I presented it above. It has two components: the first is the claim that the measure of accuracy — namely, the negative Brier score — measures cognitive value; the second is a decision-theoretic norm, namely, [Straight] Dominance. From these, via a mathematical theorem, Joyce derives Probabilism. I claim that any evidential norm can be derived in the same manner. We retain the first component, namely, the identification of cognitive value with accuracy — this is the core of veritism. But we change the decision-theoretic norm. Different decision-theoretic norms give rise to different credal norms, some of which are the evidential norms we wish to recover.”³⁵

More concisely, Pettigrew’s case for maintaining that there are scoring rule arguments that provide purely alethic justifications for the main Bayesian norms goes as follows. First, Pettigrew makes the following claim:

Claim. If one can derive a norm using a purely alethic epistemic utility function (a function that takes only considerations regarding accuracy into account) and a decision-theoretic rule that employs this utility function, then one has provided a purely alethic justification for that norm.

Second, Pettigrew holds that the utility function associated with the Brier score is purely alethic. Third, Pettigrew notes that there are scoring rule arguments which show how one can derive the main Bayesian norms by plugging this utility function into various decision-theoretic dominance rules. Together, these three claims entail that there are scoring rule arguments which provide purely alethic justifications for these Bayesian norms.

But this argument is unsound. It’s unsound because Claim is false. Deriving a norm using an alethic utility function and a decision rule doesn’t show that we have a purely alethic justification for that norm. For while the utility function may be alethic, the decision rule need not be. And if we employ a decision rule that takes factors other than accuracy into account, then our justification of the norm won’t be purely alethic.

³⁵Pettigrew (2013b), p591-592.

Let's look at a couple examples. Consider a decision rule that takes a utility function, ignores it, and then declares every option permissible. One could use the Brier-based utility function and this decision rule to derive:

Permission: Every credence function is permissible.

Given Claim, it follows that we've provided a purely alethic justification for this norm. But, of course, we've done nothing of the kind — the derivation of this norm has nothing to do with accuracy.

For another example, consider a decision rule that takes a utility function, evaluates the expected utility of each credence function (relative to the credences of the subject whose epistemic obligations we're assessing), and takes a credence function to be permissible *iff* it minimizes expected utility. One could use the Brier utility function and this decision rule to derive the negation of Probabilism. Given Claim, it follows that we have a purely alethic justification for rejecting Probabilism. Again, this is clearly incorrect — for while this decision rule and utility function combination does take accuracy into account, it's aiming away from the truth, not towards it.

Here's a third example. Let G be the proposition that God exists. And let's say that cr is G dominated by cr^* for the subject *iff* $cr(G) < cr^*(G)$. Now suppose someone adopted, as decision rules, both Straight Dominance and:

G Dominance: If a subject is such that (i) cr is G dominated by cr^* , and (ii) cr^* is not G dominated by any cr^{**} , then it is impermissible for her to adopt cr .³⁶

Using the Brier-based utility function score and these decision rules, one can derive:

Theism: It's impermissible to adopt any credence function for which $cr(G) < 1$.³⁷

Given Claim, it follows that we have a purely alethic justification for Theism. But this is wrong. For even though these decision rules do take accuracy into account, and are to some extent aiming towards the truth, accuracy isn't the only thing they're taking into account. Thus this derivation hasn't provided us with a purely alethic justification for Theism.

So Pettigrew's case for why these scoring rule arguments provide purely alethic justifications for evidential Bayesian norms is unconvincing. For his case hangs crucially on the truth of Claim — that being able to derive a norm from a purely alethic utility function and a decision-theoretic rule provides a purely alethic justification for that norm. But, as we've seen, Claim is false. For one can't determine whether we have a purely alethic justification just by ensuring that the utility function one employs is alethic. One also has to ensure that

³⁶For the reasons discussed in section 3, the second clause of both Straight Dominance and G Dominance should be “(ii) cr^* is unimpeachable”, or (if straight dominance and G dominance are the only normatively relevant factors) “(ii) cr^* is neither straight dominated nor G dominated by any cr^{**} ”. But since I'm bracketing these issues for the purposes of this section, I've just followed Pettigrew in formulating these principles using the Dominance (v3) schema.

³⁷There are, of course, plenty of historical examples of people arguing that we're rationally committed to a norm like this. But regardless of what one thinks of the merits of those arguments, no one would take *this* argument to provide a purely alethic justification for Theism. (Among other things, one can swap G with any proposition one likes — say, the proposition that goblins exist — and get an argument of exactly the same form.)

the decision-theoretic rule being employed is alethic, and that the combination of the utility function and the decision-theoretic rule remain alethic.

Let's take a step back from the particular details of Pettigrew's argument. One can see more generally that the scoring rule arguments that have been offered are, by themselves, of little use in addressing the initial worry regarding whether one can provide purely alethic justifications for evidential norms.

Recall the initial worry sketched in section 1. Purely alethic justifications can only appeal to the goal of accuracy, but evidential norms also seem to aim at the goal of having beliefs that line up with the evidence. And the goals of accuracy and lining up with the evidence can, it seems, come apart. So it's not clear how one could provide purely alethic justifications for evidential norms.

Now, the scoring rule arguments for evidential norms that have been offered all employ decision-theoretic rules in which evidence plays a key normative role. The use to these decision-theoretic rules suggests that these derivations are taking potentially non-alethic factors — evidence — into account. And we haven't been given any story for how these appeals to evidence follow from purely alethic goals. So these scoring rule arguments for evidential norms don't ultimately do anything to address the initial worry we started with. They simply shift the burden from that of showing how to provide purely alethic justifications for evidential norms to that of showing how to provide purely alethic justifications for evidence-employing decision-theoretic principles. And this burden is no easier to meet.³⁸

5 On the Possibility of Purely Alethic Justification

In the previous section we considered whether the scoring rule arguments that have been offered for evidential Bayesian norms yield purely alethic justifications for those norms. I argued that in order to show that these arguments provide purely alethic justifications, one would need to show that the decision-theoretic principles these arguments employ are themselves purely alethic. And this has yet to be done. But this discussion leaves open the possibility that someone could provide such a justification. Let's now turn to consider whether providing such a justification is possible. That is, let's consider the question of whether providing a purely alethic justification for evidential norms is possible.

In order to provide a purely alethic justification for a norm, the norm itself needs to be purely alethic; that is, the norm itself needs to make prescriptions which depend only on "truth-seeking" or alethically acceptable factors.³⁹ If changes to alethically unacceptable factors

³⁸This is not, of course, to say that this burden can't be met. (For example, one can see some of the work done by Easwaran (2014) as an attempt to discharge part of this burden.) But by and large, little has been done to try to provide purely alethic justifications for the decision-theoretic principles that the arguments sketched in section 2.3 employ.

³⁹The claim that one can only provide purely alethic justifications for purely alethic norms is, strictly speaking, only true for what one might call "epistemic ur-norms": conjunctions of every epistemic norm, which yield all of the facts regarding what's epistemically permissible (with respect to a certain mode of evaluation). To see why, suppose one provided a purely alethic justification for some principle X . One could decompose X into two principles, $X1$ and $X2$, where $X1$ imposes all and only the normative constraints that X imposes in cases involving something purple,

change the norm's prescriptions, then it won't be possible to provide a purely alethic justification for the norm — for the norm takes factors other than alethic factors into account. So if we want to determine whether it's possible to provide a purely alethic justification for a given norm, the first thing to do is to assess whether the norm's prescriptions take only alethically acceptable factors into consideration.

Of course, determining whether a factor is alethically acceptable can be tricky. For whether a factor is alethically acceptable depends on what kind of normative evaluation we're performing.

To see why, let's first consider a similar issue in ethics. Suppose you're a consequentialist. Then you'll hold that the correct moral norms will make prescriptions which only depend on "best-outcome seeking" or "consequentially acceptable" factors. But what factors count as consequentially acceptable will depend on the kind of normative evaluations you're concerned with. For example, you might be concerned with:

objective evaluations: what a subject should do given what will bring about the best results,

subjective evaluations: what a subject should do given what she believes will bring about the best results, or

prospective evaluations: what a subject should do given what she should believe will bring about the best results.

Consider your beliefs about the outcomes of acts. Whether these beliefs are a consequentially acceptable factor will depend on which of these kinds of evaluations you're performing. If you're performing an objective evaluation, then your beliefs about the outcomes of acts are independent of what you should do, and (thus) such beliefs are not a consequentially acceptable factor. If, on the other hand, you're performing a subjective evaluation, then your beliefs about the outcomes of acts are essential to determining what you should do, and thus are a consequentially acceptable factor.

If you're a Veritist, then you'll hold that the correct epistemic norms will make prescriptions which only depend on "truth seeking" or "alethically acceptable" factors. But what factors count as alethically acceptable will depend on the kind of normative evaluations you're concerned with; e.g.:

objective evaluations: what a subject should believe given what is true,

subjective evaluations: what a subject should believe given what she believes to be true, or

and X_2 imposes all and only the normative constraints that X imposes in cases not involving something purple. Then we'd have a purely alethic justifications for X_1 (since we have a purely alethic justification for X , and X entails X_1), even though X_1 is not purely alethic (because its prescriptions partially depend on non-alethic factors; e.g., whether purple things are involved).

That said, except for special cases in which the individual norms are gruesome decompositions of purely alethic norms (as in the case described above), the claim will also hold for individual norms. I.e., unless it's part of a gruesome decomposition, an individual norm will be susceptible to purely alethic justification *iff* it's purely alethic. And since the kinds of individual norms discussed in the literature are not gruesome decompositions of this kind, we can generally proceed as if this claim holds for individual norms without worry.

prospective evaluations: what a subject should believe given what she should believe to be true.

So consider what factors should play into one's verdicts if one is concerned with purely alethic justifications. If one is performing an objective evaluation, then the only factors relevant to the verdict should be the cr being evaluated, and the truth values of the propositions cr assigns values to (yielding a verdict which is relative to the world w at which we're evaluating cr). If one is performing a subjective evaluation, then the only factors relevant to the verdict should be the cr being evaluated, and the subject's credences (cr_s) in the propositions cr assigns values to (yielding a verdict which is relative to the subject for which we're evaluating cr). And if one is performing a prospective evaluation, then the only factors relevant to the verdict should be the cr being evaluated, and the credences the subject should have (cr_{sp}) in the propositions cr assigns values to (yielding a verdict which is relative to the subject for which we're evaluating cr).

What kind of evaluation are we concerned with here? At first glance, it doesn't seem like we're concerned with objective evaluations of what a subject should believe. For what a subject objectively ought to believe, it seems, is just whatever's true, and that yields a distinctly uninteresting batch of prescriptions. Likewise, at first glance, it doesn't seem like we're concerned with proscriptive evaluations of what a subject should believe. For on this approach one would be required to plug in what a subject should believe in order to figure out what a subject should believe, which threatens to be circular. This suggests that it's something like subjective evaluation that we're concerned with. Moreover, a perusal of the literature on scoring rule arguments makes it clear that it's subjective evaluation that this literature is concerned with. (E.g., consider the focus on the *expected* epistemic utility of adopting a credence function, the appeal to constraints like Propriety which presuppose the relevance of maximizing expected utility, and so on.) So in what follows, let's narrow our focus to the following question: is providing a purely alethic *subjective* justification for evidential norms possible?

In order to provide a purely alethic justification for a norm, the norm must make prescriptions which depend only on alethically acceptable factors. Given a subjective mode of evaluation, the only alethically acceptable factors are the cr being evaluated and the subject's credences cr_s . So the following is a necessary condition on being able to provide a purely alethic subjective justification for a norm:

Supervenience: One can't alter the verdicts of the norm without altering either the credence function cr being evaluated or the credences of the subject cr_s with respect to which the evaluation is being performed.

At first glance, none of the evidential Bayesian norms we considered in section 2.3 appear to satisfy this condition. In all three cases, it seems, one can flip a verdict from permissible to impermissible by tweaking the subject's evidence, while leaving cr and cr_s unchanged.

Consider Conditionalization. Suppose a subject with a credence of $1/2$ in A receives A as evidence, and the credence function cr being evaluated assigns a credence of 1 to A . According to Conditionalization, adopting cr could be permissible. Suppose, on the other hand, that a subject with a credence of $1/2$ in A receives only the tautology \top as evidence, and the credence

function cr being evaluated assigns a credence of 1 to A . Then Conditionalization would hold that adopting cr is impermissible.

Or consider the Principal Principle. Suppose a subject's total evidence is that the chance of A is $1/2$, and the credence function cr being evaluated assigns a credence of $1/2$ to A . Then according to the Principal Principle, adopting cr could be permissible. On the other hand, suppose the subject's total evidence is that the chance of A is 1, and the credence function cr being evaluated assigns a credence of $1/2$ to A . Then the Principal Principle would hold that adopting cr is impermissible.

Or consider the Indifference Principle. Suppose a subject's total evidence is w_1 , and the credence function cr being evaluated is such that $cr(w_1) \neq cr(w_2)$. Then according to the Indifference Principle, adopting cr could be permissible. On the other hand, suppose a subject's total evidence is the tautology \top , and the credence function cr being evaluated is such that $cr(w_1) \neq cr(w_2)$. Then the Indifference Principle would maintain that adopting cr is impermissible.

So, it seems, none of the evidential Bayesian norms we've considered satisfy Supervenience. More generally, since all we need to construct a counterexample to the condition is for the norm's verdicts to take the subject's evidence into account, it seems that no evidential norm could satisfy this condition. Thus, it seems, it's not possible to provide a purely alethic subjective justification for evidential norms.

There is, however, a way to avoid this conclusion. Howson & Urbach (1993) have suggested a picture of evidence on which a subject's evidence is just whatever the subject has a credence of 1 in. If we adopt this picture of evidence, then one can't change a subject's evidence without also changing her credences. Thus on this picture of evidence all of the counterexamples to Supervenience considered above, which involve changing the subject's evidence while leaving her credences the same, are impossible. So if we adopt something like Howson and Urbach's picture of evidence, we can maintain that evidential norms *can* satisfy Supervenience.⁴⁰

⁴⁰Note that to get the evidential norms to satisfy Supervenience, it's crucial that one's account of evidence be such that one has E as evidence *iff* one has a credence of 1 in E , as per Howson and Urbach's account. Merely having the entailment go one way or the other won't suffice.

To see this, first consider accounts on which having E as evidence entails having a credence of 1 in E , but not vice versa. E.g., consider something like Williamson's (2000) account, on which one has evidence E *iff* one knows E . Suppose a subject comes to have a credence of 1 in E by coming to know E . Then, given this account of evidence, Conditionalization will prescribe conditionalizing on E . On the other hand, suppose a subject comes to have a credence of 1 in E but doesn't know E . Then Conditionalization won't prescribe conditionalizing on E . So given this account of evidence, subjects with the same credences can be given different prescriptions by evidential norms. So evidential norms will fail to satisfy Supervenience.

Second, consider accounts on which having a credence of 1 in E entails that E is evidence, but not vice versa. E.g., consider an account on which one has evidence E *iff* one has a credence of 1 in E *or* one has been told E by an authoritative source. Suppose a subject has a credence of $1/2$ in E and has been told E by an authoritative source. Then, given this account of evidence, Conditionalization will prescribe conditionalizing on E . On the other hand, suppose a subject has a credence of $1/2$ in E and hasn't been told E by an authoritative source. Then Conditionalization won't prescribe conditionalizing on E . So given this account of evidence, subjects with the same credences can be given different prescriptions by evidential norms. So again, evidential norms will fail to satisfy Supervenience.

But defending the possibility of purely alethic justifications for evidential norms this way has some big drawbacks. Many people (like me) don't like Howson and Urbach's thin picture of evidence, and would prefer to endorse a more substantive picture of evidence.⁴¹ And this way of defending the possibility of purely alethic justifications for evidential norms is of little help to such people.

Moreover, proponents of substantive evidence will be actively hostile to the prospect of justifying versions of these evidential norms that employ Howson and Urbach's notion of evidence. This is because they'll take the Howson and Urbach versions of these norms to be false, and thus impossible to successfully justify.

Consider Conditionalization. Suppose a subject is hit on the head with a rock, and as a result of this mental trauma, comes to have a credence of 1 that unicorns exist.⁴² The version of Conditionalization that employs Howson and Urbach's notion of evidence will take this to be a case of receiving evidence, and will maintain that the subject should adopt the credences they'd get by conditionalizing on the proposition that unicorns exist. Pretty much every version of Conditionalization that employs a more substantive notion of evidence, on the other hand, will not take this to be a case of receiving evidence, and will maintain that the subject's credences should remain the same. So these two different understandings of Conditionalization will yield conflicting prescriptions. And thus proponents of the substantive evidence version of Conditionalization will take the Howson and Urbach version to be false, and hence something which can't be successfully justified.

The same is true for other evidential norms, like the Principal Principle and the Indifference Principle. The Howson and Urbach versions of these norms will yield prescriptions that conflict with those of the substantive evidence versions of these principles. And thus proponents of the substantive evidence versions of these principles will take the Howson and Urbach versions of these principles to be false, and any purported justifications of them erroneous.

So proponents of substantive evidence will not just fail to be mollified by the possibility of providing a purely alethic justification for the Howson and Urbach versions of these norms. They'll be actively hostile to the idea of there being successful justifications for these norms (whether purely alethic or not), because they'll take these norms to be false.

So is it possible to provide purely alethic (subjective) justifications for evidential norms? Unless one adopts something like Howson and Urbach's picture of evidence, and the corresponding understanding of these norms, the answer is "no". And, as proponents of substantive evidence have argued, there are a number of reasons to dislike Howson and Urbach's picture of evidence.^{43,44}

⁴¹I take Howson and Urbach's picture of evidence to be "thin" in the sense that it provides a reductive analysis of evidence in terms of other components of the Bayesian apparatus (i.e., credence). Proponents of more substantive pictures of evidence, which can't be reduced in this way, include Lewis (1996), Williamson (2000), Neta (2008) and Dunn (2010).

⁴²For a lengthier discussion of this kind of case, and of different ways of thinking about evidence, see Dunn (2010).

⁴³For an in-depth discussion of these issues, and further references, see Dunn (2010).

⁴⁴It's natural to wonder here whether the assumption that evidence is veridical could provide one with a way to construct purely alethic justifications for evidential norms. If we were concerned with objective justification, and were assessing what a subject should believe given what's true, then the assumption that any proposition one receives

6 Purely Alethic Justification and Content-Sensitive Norms

In the previous section, we considered whether it's possible to provide purely alethic (subjective) justifications for evidential norms. I argued that it's not, unless you buy something like the contentious picture of evidence proposed by Howson and Urbach, and the corresponding contentious understanding of these evidential norms. In this section let's consider how things go if we do grant something like Howson and Urbach's picture of evidence. In particular, let's consider whether, even given this picture of evidence, it's possible to provide purely alethic justifications for evidential norms whose prescriptions concern particular kinds of propositions, like the Principal Principle.

Here is a *prima facie* reason to worry about the possibility of such justifications. One can only provide a purely alethic justification for a norm that's purely alethic — a norm that cares only about accuracy, and directs subjects to adopt the most accurate beliefs they can.⁴⁵ In the case of purely alethic subjective justification, this will be accuracy as assessed by the evaluating credence function. Every other factor about a subject's doxastic state should be irrelevant. For example, the *content* of the beliefs we're evaluating shouldn't matter, just whether they're accurate. Now, it seems like norms whose prescriptions concern particular kinds of propositions will necessarily take the content of the propositions we're evaluating into account. For example, the Principal Principle will treat some kinds of propositions — propositions about chances — differently from other propositions. But if this is right, then it seems norms whose prescriptions concern particular kinds of propositions, like the Principal Principle, can't be given purely alethic justifications.

Let's see how one might make this worry more precise. As we saw in the last section, one can't provide a purely alethic justification for a norm if its prescriptions depend on alethically unacceptable factors. Given a subjective mode of evaluation, this yields Supervenience as a necessary condition on being able to provide a purely alethic justification for a norm. But this isn't the only constraint that a norm must satisfy in order to be susceptible to purely alethic justification. For a norm could make prescriptions that depend only on alethically acceptable factors (and thus satisfy Supervenience), and yet use these factors to direct subjects away from the truth, or in directions orthogonal to the truth.

as evidence is true might help. But if we're concerned with subjective justification, and assessing what a subject should believe given what she believes to be true, then it's whether the subject *believes* her evidence is true, not whether it's actually true, that will matter. Thus while this veridicality assumption might help with providing purely alethic objective justifications for evidential norms, it won't help to provide purely alethic subjective justifications for evidential norms. And it's subjective justifications that the scoring rule literature we're discussing is concerned with.

(Note that Howson and Urbach's account of evidence entails that subjects always believe their evidence is true. So insofar as the thought was to improve our odds of being able to provide purely alethic justifications for evidential norms by adopting an account of evidence on which evidence is veridical, or has whatever analog of veridicality that's relevant to the mode of evaluation we're concerned with, then Howson and Urbach's account already has the desired features in the context of subjective evaluations.)

⁴⁵In the case of purely alethic *subjective* justification, this will be accuracy assessed with respect to the evaluating credence function. Again, strictly speaking, this is only true for the epistemic "ur-norm", not individual norms; see footnote 39. But, again, given the kinds of individual norms we're concerned with, we can put this caveat aside.

Some further constraints on such norms are suggested by the discussion of alethic utility functions in section 2.2. We saw it suggested there that alethic utility functions should satisfy constraints like Truth-Directedness and Extensionality. One natural thought is to try to transform these constraints on utility functions into constraints on norms. Let's see how one might do that. In particular, given the *prima facie* worry sketched above, let's see how one might transform the Extensionality constraint on utility functions into an analogous constraint on norms.⁴⁶

Recall the form of the Extensionality constraint on utility functions:

Extensionality: Let $f_{w,w^*} : F \rightarrow F$ be a bijection between true and false propositions at w with true and false propositions at w^* ; i.e. if $f_{w,w^*}(A) = B$, then either A is true at w and B is true at w^* , or A is false at w and B is false at w^* . If cr and cr^* are such that, for all A , $cr(A) = cr^*(f_{w,w^*}(A))$, then $U(cr, w) = U(cr^*, w^*)$.

This condition requires utility functions to not care about the content of propositions. It says that if you have a mapping f from propositions at w to propositions at w^* that preserves truth value, and if the values cr assigns to these propositions are the same as the values cr^* assigns to what f maps them to, then the epistemic utility of holding cr at w is the same as the epistemic utility of holding cr^* at w^* . Why? Because (the proponent of Extensionality says) all that matters to the epistemic utility of a credal state is how close or far its beliefs are from what's true. Thus permuting *which* beliefs are close or far from the truth shouldn't make a difference to the utility of holding those beliefs.

How should we formulate an analogous constraint on norms, given that we're interested in assessing credal states using a subjective mode of evaluation? To get such a constraint, we need to change a couple things. First, since we're performing a subjective evaluation, we want a constraint which employs what the evaluating subject *believes* to be true or false in place of what's actually true or false. So we'll want to replace appeals to w/w^* with appeals to evaluating credence functions cr_s/cr_s^* . And we'll want to replace appeals to the truth values of propositions at w/w^* with appeals to the credence assigned to those propositions by cr_s/cr_s^* . Second, since we want a constraint on prescriptions instead of utility assignments, we want a constraint which bears on the permissibility of credal states instead of the utility assignments to credal states. So we'll want to replace the requirement that $U(cr, w) = U(cr^*, w^*)$ with requirement that cr have the same deontic status with respect to cr_s as cr^* has with respect to cr_s^* .

Making these substitutions gives us:

Extensional Prescriptions: Let $f_{cr_s, cr_s^*} : F \rightarrow F$ be a bijection that pairs propositions assigned a given value by cr_s with propositions assigned the same value by cr_s^* ; i.e., if $f_{cr_s, cr_s^*}(A) = B$, then $cr_s(A) = cr_s^*(B)$. If cr and cr^* are such that, for all A , $cr(A) = cr^*(f_{cr_s, cr_s^*}(A))$, then cr has the same deontic status with respect to cr_s as cr^* has with respect to cr_s^* .

⁴⁶To be clear, I won't be providing anything like a theorem linking the adoption of Extensionality to the adoption of an analogous constraint on norms (which I call Extensional Prescriptions, below). Rather, I'll be arguing that (i) Extensional Prescriptions is the analog of Extensionality in the case of norms, and (ii) the same thoughts that motivate Extensionality also motivate Extensional Prescriptions.

This condition requires *norms* to not care about the content of propositions. It says that if you have a mapping f from propositions to propositions such that the values cr_s assigns to these propositions are the same as the values cr_s^* assigns to what f maps them to, and if the values cr assigns to these propositions are the same as the values cr^* assigns to what f maps them to, then the deontic status of holding cr (assessed with respect to cr_s) is the same as the deontic status of holding cr^* (assessed with respect to cr_s^*). Why? Because (the proponent of Extensional Prescriptions says) given a subjective mode of evaluation, all that matters to the deontic status of a credal state is how close or far it is from what the evaluating credal state believes to be true. Thus permuting *which* beliefs are close or far from what the evaluating credal state believes to be true shouldn't make a difference to the deontic status of the credal state.⁴⁷

Compatibility with Extensional Prescriptions, like compatibility with Supervenience, is a plausible necessary condition on being able to provide a purely alethic subjective justification for a norm. Thus if content-sensitive norms like the Principal Principle are incompatible with Extensional Prescriptions, then it seems that such norms can't be given purely alethic subjective justifications. And given some widely held assumptions, content-sensitive norms like the Principal Principle *will* conflict with Extensional Prescriptions. Let's see why this is so in the case of the Principal Principle, by working through an example.

Let cr_s be a credence function such that:

$$\begin{aligned} cr_s(\langle ch(F) = 1/4 \rangle) &= 2/3 & cr_s(F) &= 1/3 \\ cr_s(\langle ch(F) = 1/2 \rangle) &= 1/3 & cr_s(G) &= 1/2 \end{aligned}$$

Note that these assignments are compatible with the Principal Principle.⁴⁸ Let's stipulate that all of cr_s 's other assignments are compatible with the Principal Principle as well. Let's further stipulate that cr_s 's assignments are compatible with every other constraint on credences that one wants to impose, as long as these constraints are compatible with the Principal Principle.⁴⁹

Even though cr_s satisfies all these constraints, one might still hold that cr_s is impermissible

⁴⁷In the text I grant Extensionality, and then proceed, via the most straightforward way of transforming Extensionality into a constraint on norms, to obtain Extensional Prescriptions. But I think Extensional Prescriptions is not the most plausible constraint on norms of its kind. A more plausible principle would further restrict the bijections it appeals to, in order to avoid potential worries regarding the equivalence of probability-satisfying and probability-violating credence assignments. Such a principle, call it Extensional Prescriptions*, would add the following requirement to the bijection f : it must map worlds to worlds, and map propositions corresponding to each set of worlds (e.g., $\{w_1, w_2, \dots\}$) to propositions corresponding to the set of worlds those worlds are mapped to (i.e., $\{f(w_1), f(w_2), \dots\}$).

Working backwards, one could appeal to Extensional Prescriptions* to argue that we should adopt "Extensionality*" (a principle like Extensionality, but with the further constraint on the bijections f) instead of Extensionality. For Extensionality* is the most straightforward transformation of Extensional Prescriptions* into a constraint on utility functions.

That said, these issues have little bearing on the topic of this section. For content-sensitive norms like the Principal Principle will be incompatible with both Extensional Prescriptions and Extensional Prescriptions*, so it won't matter which principle we adopt. For this reason, I'll put these complications aside, and proceed in the text as if Extensional Prescriptions was unproblematic.

⁴⁸Since $cr_s(F) = \sum_{x \in [0,1]} cr_s(\langle ch(F) = x \rangle) \cdot x$, as the Principal Principle requires (see Lewis (1980)).

⁴⁹Thus these constraints can't include the Indifference Principle described in section 2.3, since this principle conflicts with the Principal Principle (c.f. footnote 32).

if evaluated with respect to a credence function that assigns very different values than cr_s . For such a credence function will take cr_s 's assignments to be far from the truth. But everyone in this debate who accepts the Principal Principle will grant that cr_s is permissible with respect to itself.

Premise 1. cr_s is permissible with respect to cr_s .

Let cr_s^* be a credence function which flips cr_s 's assignments to F and G , and assigns the same values as cr_s to every other proposition. Thus:

$$\begin{aligned} cr_s^*(\langle ch(F) = 1/4 \rangle) &= 2/3 & cr_s^*(F) &= 1/2 \\ cr_s^*(\langle ch(F) = 1/2 \rangle) &= 1/3 & cr_s^*(G) &= 1/3 \end{aligned}$$

Note that these assignments are incompatible with the Principal Principle.⁵⁰ So, given the Principal Principle, it will always be impermissible to adopt cr_s^* , even when evaluated with respect to itself.

Premise 2. cr_s^* is impermissible with respect to cr_s^* .

Now consider a bijection f that maps F to G , G to F , and every other proposition to itself.⁵¹ This bijection will be such that, for all A , $cr_s(A) = cr_s^*(f(A))$. Thus Extensional Prescriptions will require that if cr_s is permissible with respect to cr_s , then cr_s^* is permissible with respect to cr_s^* .

Premise 3. cr_s is permissible with respect to cr_s iff cr_s^* is permissible with respect to cr_s^* .

Thus given the Principal Principle and Extensional Prescriptions we get a contradiction: cr_s is permissible (with respect to itself), cr_s^* is impermissible (with respect to itself), but both must have the same deontic status (with respect to themselves). So if we adopt the Principal Principle, we must reject Extensional Prescriptions. And if we reject Extensional Prescriptions, we have to give up on the thought that we can provide purely alethic justifications for every prescription. Thus one can't both adopt the Principal Principle and hold that we can provide purely alethic justifications for every prescription.

More generally, we can see that principles which make prescriptions regarding particular kinds of propositions (like propositions about chances) will conflict with Extensional Prescriptions. For Extensional Prescriptions says that if a credal state is permissible, then any permutation of its credences that keeps their distance from the evaluating credence function the same should also be permissible. But, for principles which are content-sensitive, such permutations can move us from credal states that satisfy the principle to credal states that violate it. So content-sensitive principles aren't susceptible to purely alethic justification.

Let's conclude by considering a reply to this worry, suggested by some remarks by Pettigrew (2013b) and Joyce (2016). To provide a purely alethic justification for a norm, one needs to provide a justification that doesn't appeal to anything besides the goal of having accurate beliefs. But one might try to escape the worries sketched above by understanding "accuracy"

⁵⁰Since $1/2 = cr_s^*(F) \neq \sum_{x \in [0,1]} cr_s^*(\langle ch(F) = x \rangle) \cdot x = 1/3$.

⁵¹If we adopt Extensional Prescriptions* instead of Extensional Prescriptions, as I suggest in footnote 47, then we'll generally need to spell out a more extensive bijection. But this is relatively straightforward to do.

in a looser way than I've been understanding it here. Pettigrew (2013b) suggests a notion of accuracy according to which a credal state is perfectly accurate when it lines up with the chances, not when it lines up with the truth. And Joyce (2016) endorses a similar move, arguing that the Principal Principle is compatible with Veritism by maintaining that "proponents of accuracy-centered approaches will see [lining up one's credences with the chances] as comprising an essential part of the duty to rationally pursue accuracy".⁵² Given one of these modified notions of accuracy, one can construct an accuracy-based argument for the Principal Principle. And while these accuracy-based justifications for the Principal Principle will yield a norm that conflicts with Extensional Prescriptions, one might argue that this only goes to show that Extensional Prescriptions is too stringent a constraint, given this more relaxed understanding of accuracy.

The trouble with this move is that it's a bait and switch. Recall what we were initially concerned with: we were concerned with the notion of accuracy (and the corresponding notion of purely alethic justification) relevant to Veritism. And the notion of accuracy relevant to Veritism was closeness to the truth. Now, one can modify one's notion of accuracy so that it tracks something other than closeness to the truth. But one can't, at the same time, claim that this is the notion of accuracy relevant to Veritism.⁵³

Here's another way to put the point. Call this modified notion of accuracy "accuracy*", call a justification which doesn't appeal to anything other than the goal of having accurate* beliefs "purely alethic justification*", and call the thesis that all epistemic norms can be given purely alethic justifications* "Veritism*". In response to the worries sketched above, this reply notes that one can provide purely alethic justifications* for evidential norms like the Principal Principle, and thus can reconcile evidential norms with Veritism*. But this reply doesn't address these worries, it just changes the topic. For these worries concerned whether one could reconcile evidential norms with Veritism, not whether one could reconcile evidential norms with Veritism*.

Another worry for this move is that it seems one could make the same kind of maneuver to reconcile Veritism with virtually any norm. In the maneuver suggested by Pettigrew and Joyce, one modifies the notion of accuracy so that it assesses closeness to the chances instead of closeness to the truth, in order to provide a purely alethic justification for the Principal Principle. But one could equally well modify the notion of accuracy so that it assesses closeness to Donald Trump's credal state instead of closeness to the truth, in order to provide a purely alethic justification for Trumpism, the claim that one should line up one's beliefs with Trump's.⁵⁴ Or one could modify the notion of accuracy so that it assesses closeness to one's current credal state instead of closeness to the truth, in order to provide a purely alethic justification for Calcification, the claim that one should never change one's beliefs.⁵⁵

⁵²See Joyce (2016), p16.

⁵³To press the point further, recall that the intuitive motivation for Veritism was precisely that one shouldn't take any factors other than closeness to the truth into account. Yet these modified notions of accuracy do just that.

⁵⁴Or should line up with what one believes Trump's beliefs to be (if we want to precisely mirror the form of the Principal Principle).

⁵⁵Or should line up with what one believes one's beliefs to be (if we want to precisely mirror the form of the

Now, perhaps the thought is that this chance-tracking notion of accuracy is tenable, while these other modified notions of accuracy are not, because the chances (unlike Trump's credal state) are closely tied to the truth. First, if this is the thought, then it's something that needs to be demonstrated, not assumed. And demonstrating how the chances are tied to the truth in the right way is precisely what we needed to do, in order to provide a purely alethic justification for the Principal Principle, when we employed our *original* notion of accuracy. So, at best, this version of the maneuver just brings us back to where we started.

Second, if the literature on these issues is any guide, then providing a non-question begging argument for why lining up one's beliefs with the chances is a good way to get true beliefs is a project with poor prospects. (E.g., see Strevens (1999) for an extended criticism of such arguments.) And even if we put that literature aside, we've just seen a positive reason for doubting that providing a rationale for the Principal Principle along these lines is possible. For if it were possible, then one should be able to provide a purely alethic justification for the Principal Principle in a way that's compatible with constraints like Extensional Prescriptions. And as we've seen, this can't be done.

In the last section we saw that it's not possible to provide a purely alethic justification for evidential norms unless one adopts something like Howson and Urbach's versions of these norms. In this section we saw that even if we do adopt the Howson and Urbach's versions of these norms, then, if we understand "purely alethic justification" in the way that's relevant to the issues at hand, there will be evidential norms for which it's not possible to provide purely alethic justifications. For norms which are susceptible to purely alethic justification should be compatible with constraints like Extensional Prescriptions. And evidential norms which are content-sensitive, like the Principal Principle, are not.

7 Conclusion

For most of this paper I've been taking Joyce and Pettigrew to be advocating the All-Accuracy Account. But some passages, such as this passage from Joyce, suggest a much weaker claim:

"Proponents of accuracy-centered epistemology will want to... [maintain] that all facts about justification are interpreted as facts about the rational pursuit of doxastic accuracy. This would make it a truism (on the level of 'evidence for X is evidence for X 's truth') that believers are justified in holding credences exactly to the extent that their evidence makes it reasonable for them to expect those credences to be accurate (in the aggregate)."⁵⁶

This passage suggests that, perhaps, Joyce is only claiming this: for any norm, one can construct an argument for that norm that has the form of a scoring rule argument. This is a claim that I think virtually everyone would agree with. After all, similar claims have been made in a number of places (e.g., see Stalnaker (2002)), and have been allowed to pass without challenge. And this claim seems clearly true: given any choice of utility functions, and any choice

Principal Principle). The name "Calcification" comes from Christensen (1991).

⁵⁶Joyce (2016), p18.

of decision rules, one *can* construct an argument for any set of prescriptions one wants. Indeed, we don't even need this much — merely granting one free choice of decision rules is enough.⁵⁷

If this is the claim that Joyce has in mind, then none of the arguments offered in this paper are a threat to it. But this claim is uninteresting. For it doesn't bear on any substantive issues regarding justification, epistemic rationality, or epistemic norms.

In any case, questions regarding views like the All-Accuracy Account are interesting, regardless of who adopts them. This paper has offered a critical look at such views. In light of this, one might conclude that views like the All-Accuracy Account are irreparably flawed, and should be rejected. Or one might just conclude that there are a number of issues that proponents of views like the All-Accuracy Account need to address. Either way, the tension between Veritism and evidential norms has yet to be resolved.⁵⁸

References

- Christensen, D. (1991). Clever Bookies and Coherent Beliefs. *Philosophical Review*, 100(2), 229–247.
- Dunn, J. (2010). *Bayesian Epistemology and Having Evidence*. Ph.D. thesis, University of Massachusetts, Amherst.
- Easwaran, K. (2013). Expected Accuracy Supports Conditionalization—and Conglomerability and Reflection. *Philosophy of Science*, 80(1), 119–142.
- Easwaran, K. (2014). Decision theory without representation theorems. *Philosophers' Imprint*, 14(27).
- Easwaran, K., & Fitelson, B. (2012). An 'Evidentialist' Worry About Joyce's Argument for Probabilism. *Dialectica*, 66(3), 425–433.
- Goldman, A. (2002). The Unity of the Epistemic Virtues. In *Pathways to Knowledge*, (pp. 51–72). New York: Oxford University Press.
- Greaves, H., & Wallace, D. (2006). Justifying conditionalization: Conditionalization maximizes expected epistemic utility. *Mind*, 115(459), 607–632.

⁵⁷To see this, note that given any utility function, one can construct a decision rule which effectively ignores that utility function and yields the prescriptions one would get by applying decision rule X to utility function Y (for any X and Y). (How can the rule figure out these prescriptions if they rely on the assignments of utility function Y , and the rule has some other utility function $Z \neq Y$ as an argument? Well, note that a free choice of decision rules entails a free choice of what the arguments of the decision rule are. So one can construct a decision rule whose arguments include whatever one needs to deduce what Y 's assignments would be.)

⁵⁸I'd like to thank Jeff Dunn, Kenny Easwaran, Maya Eddon, and Richard Pettigrew for helpful comments and discussion.

- Hedden, B. (2015). *Reasons Without Persons: Rationality, Identity, and Time*. Oxford University Press Uk.
- Howson, C., & Urbach, P. (1993). *Scientific Reasoning: The Bayesian Approach*. Open Court.
- Joyce, J. (2009). Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief. In Franz Huber, & Christoph Schmidt Petri (Eds.) *Degrees of Belief*, vol. 342, (pp. 263–297). Synthese.
- Joyce, J. M. (1998). A nonpragmatic vindication of probabilism. *Philosophy of Science*, 65(4), 575–603.
- Joyce, J. M. (2016). Why evidentialists need not worry about the accuracy argument for probabilism.
- Leitgeb, H., & Pettigrew, R. (2010a). An Objective Justification of Bayesianism II: The Consequences of Minimizing Inaccuracy. *Philosophy of Science*, 77(2), 236–272.
- Leitgeb, H., & Pettigrew, R. (2010b). An Objective Justification of Bayesianism I: Measuring Inaccuracy. *Philosophy of Science*, 77(2), 201–235.
- Levinstein, B. A. (2012). Leitgeb and Pettigrew on Accuracy and Updating. *Philosophy of Science*, 79(3), 413–424.
- Lewis, D. (1980). A subjectivist's guide to objective chance. In R. C. Jeffrey (Ed.) *Studies in Inductive Logic and Probability*, vol. 2, (pp. 83–132). University of California Press.
- Lewis, D. (1996). Elusive Knowledge. *Australasian Journal of Philosophy*, 74(4), 549–567.
- Meacham, C. J. G. (2010). Unravelling the tangled web: Continuity, internalism, non-uniqueness and self-locating beliefs. In T. S. Gendler, & J. Hawthorne (Eds.) *Oxford Studies in Epistemology, Volume 3*, (p. 86). Oxford University Press.
- Meacham, C. J. G. (2016). Understanding Conditionalization. *Canadian Journal of Philosophy*, 45(5-6), 767–797.
- Moss, S. (forthcoming). Time-slice epistemology and action under indeterminacy. In J. Hawthorne, & T. Gendler (Eds.) *Oxford Studies in Epistemology 5*.
- Neta, R. (2008). What Evidence Do You Have? *British Journal for the Philosophy of Science*, 59(1), 89–119.
- Nozick, R. (1974). *Anarchy, State and Utopia*. Basic Books.
- Pettigrew, R. (2012). Accuracy, Chance, and the Principal Principle. *Philosophical Review*, 121(2), 241–275.

- Pettigrew, R. (2013a). A New Epistemic Utility Argument for the Principal Principle. *Episteme*, 10(1), 19–35.
- Pettigrew, R. (2013b). Accuracy and Evidence. *Dialectica*, 67(4), 579–596.
- Pettigrew, R. (2013c). Epistemic Utility and Norms for Credences. *Philosophy Compass*, 8(10), 897–908.
- Pettigrew, R. (2014). Accuracy, Risk, and the Principle of Indifference. *Philosophy and Phenomenological Research*, 91(1), n/a–n/a.
- Schoenfield, M. (2016). Conditionalization does not (in general) maximize expected accuracy.
- Stalnaker, R. (2002). Epistemic Consequentialism: Robert Stalnaker. *Aristotelian Society Supplementary Volume*, 76(1), 153–168.
- Strevens, M. (1999). Objective Probability as a Guide to the World. *Philosophical Studies*, 95(3), 243–275.
- van Fraassen, B. C. (1995). Belief and the problem of Ulysses and the Sirens. *Philosophical Studies*, 77(1), 7–37.
- Williamson, T. (2000). *Knowledge and its Limits*. Oxford University Press.