

Accuracy, Probabilism, and the Insufficiency of the Alethic

Corey Dethier

[Pre-print; forthcoming in *Philosophical Studies*]

Abstract

The best and most popular argument for probabilism is the accuracy-dominance argument, which purports to show that alethic considerations alone support the view that an agent's degrees of belief should always obey the axioms of probability. I argue that extant versions of the accuracy-dominance argument face a problem. In order for the mathematics of the argument to function as advertised, we must assume that every omniscient credence function is classically consistent; there can be no worlds in the set of dominance-relevant worlds that obey some other logic. This restriction cannot be motivated on alethic grounds unless we're also willing to accept that rationality requires belief in every metaphysical necessity, as the distinction between a priori logical necessities and a posteriori metaphysical ones is not an alethic distinction. To justify the restriction to classically consistent worlds, non-alethic motivation is required. And thus, if there is a version of the accuracy-dominance argument in support of probabilism, it isn't one that is grounded in alethic considerations alone.

0 Introduction

The best and most popular argument for probabilism is the accuracy-dominance argument, which purports to show that alethic considerations—that is, considerations relating to modalities of truth—alone support the view that an agent's degrees of belief should always obey the axioms of probability. In slightly more detail: if we assume that rationality requires us to be as close to omniscience as possible, then no

matter what informational state we're in, the best credence function or functions will be probability functions. In this paper, I argue that the argument requires supplementation: alethic considerations alone are insufficient to motivate the standard version of probabilism.

The problem is as follows. In order for the mathematics of the accuracy-dominance argument to function as advertised, we must assume that every omniscient credence function is classically consistent; there can be no worlds in the set of “dominance-relevant worlds” (Pettigrew 2021) that obey some other logic. This restriction cannot be motivated on alethic grounds unless we're also willing to accept that rationality requires belief in every metaphysical necessity, as the distinction between a priori logical necessities and a posteriori metaphysical ones is not an alethic distinction. To justify the restriction to classically consistent worlds, non-alethic motivation is required. And thus, if there is a successful version of the accuracy-dominance argument in support of probabilism, it isn't one that is grounded in alethic considerations alone.

To be clear, my intent here is not to show that probabilism is false, or even to show that the accuracy-based argument for it fails. Instead, my aim is to attack a particular interpretation of what it is that the accuracy-dominance argument shows. On the interpretation in question, the only (fundamental) epistemic value employed in the argument is the *alethic* value of accuracy to the truth.¹ Since probabilistic credence functions are closer to the truth “no matter what,” any non-probabilistic credence function is irrational in virtue of conflicting with this fundamental alethic value. My contention is that to be successful, the accuracy-dominance argument requires (implicit or explicit) appeal to additional, non-alethic, considerations—particularly considerations of justification or evidence—because the desire for accuracy alone is insufficient to motivate a distinction between logical and metaphysical necessities, and the argument needs this distinction to succeed.

The layout of the paper is simple. Section 1 sketches the accuracy-dominance argument for probabilism. Section 2 shows that the argument relies on restricting the set of possibilities to those that are classically consistent. Section 3 argues that alethic considerations are insufficient to justify the restriction. Section 4 show how introducing evidential considerations in addition to alethic ones can resolve the problem. Section 5 discusses the consequences of this move.

One final note. As already indicated, by “alethic” considerations I mean considerations relating to metaphysical modalities; by “evidential” considerations I mean

¹So far as I can tell, this is the interpretation offered by both Joyce (1998) and Pettigrew (2016). The former, for instance, says describes the argument as showing that probabilism “contributes to the basic epistemic goal” of “accurately represent[ing] the world” (Joyce 1998, 576).

considerations relating to how evidence constrains (rational) belief. As one reviewer rightly stressed, however, these aren't the only ways to use these notions. For instance, the modal distinction between epistemic mights and musts is at least arguably alethic. As we'll see in the final section, if the defender of the accuracy-dominance argument wants to appeal to alethic considerations in this (broader) sense, I have no objections—but I'll argue that doing so requires rethinking what it is that the argument shows.

1 The accuracy-dominance argument for probabilism

In its simplest form, which is all we need here, the accuracy-dominance argument for probabilism proceeds as follows:

(DA) DOMINANCE AVOIDANCE

An agent is rational only if they are not accuracy dominated.²

(CD) CREDAL DOMINANCE

An agent's degrees of belief obey the axioms of probability if and only if they are not accuracy dominated.

∴ (PR) PROBABILISM

An agent is rational only if their degrees of belief obey the axioms of probability.³

My concerns with this argument are best understood as concerns with the first premise. To see why I take the first premise to be problematic, however, we'll need to understand what motivates the second premise, namely (CD). In this section, I sketch a largely intuitive outline of the proof of what I'm calling credal dominance (those

²As Carr (2017, 2021) stresses, it may be better to think of the accuracy-dominance argument as applying to the rationality of credences or credal states rather than agents. I'm foregoing such complications here for ease of presentation.

³For the purposes of the first four sections of this essay, I'll limit "probabilism" to this claim, though there are other theses that go by the same name. For instance: Pettigrew (2021) defends what he calls "personal probabilism," according to which rationality demands a different form of coherence determined by what possibilities are open to an agent. My arguments here should not be taken to undermine that view or others like it; I'm concerned only with versions of probabilism that require logical omniscience.

interested in the technical details should see Joyce 1998, 2009; Pettigrew 2015, 2016). My objection is raised in the next section. Anyone familiar with the technical details can safely skip the rest of this section—though my terminology differs at times from that of both Joyce and Pettigrew, I’m relatively confident that any serious differences will be clear in context.

Begin with a set of propositions \mathcal{F} , which are assumed to form a Boolean algebra; they’re closed under negation and countable disjunction. A *credence function* on \mathcal{F} is a function $c: \mathcal{F} \rightarrow [0, 1]$. A *probability function* on \mathcal{F} is a credence function that is *normalized* ($c(\perp) = 0$ and $c(\top) = 1$) and (countably) *additive* (for any partition \mathbb{X} , $c(\mathbb{X}) = \sum_{X \in \mathbb{X}} c(X)$). The set of probability functions we’ll represent with $\mathcal{P}_{\mathcal{F}}$. To measure the distances between credence functions, we’ll use the “Brier score”:

$$\mathcal{B}(c', c) = \sum_{X \in \mathcal{F}} |c'(X) - c(X)|^2$$

where the lower the score is the better. To measure accuracy in particular, we’ll use the distance from the omniscient credence function at a world w , which we’ll call o_w . The crucial property that we need can be defined as follows:

ACCURACY DOMINATED

A credence function c is *accuracy dominated* if and only if there is a credence function c^* such that c^* is closer to every omniscient credence function than c .

This definition can be stated more formally as

$$\mathbf{AD}(c) := \exists c^* \supset \forall w \in \mathcal{W}_{\mathcal{F}}, \mathcal{B}(o_w, c^*) < \mathcal{B}(o_w, c)$$

Here $\mathcal{W}_{\mathcal{F}}$ is the set of possible worlds, where these are understood as the set of maximally consistent propositions in \mathcal{F} . As indicated above, what we want to prove is that a credence function is accuracy dominated if and only if it is not a probability function: $c \notin \mathcal{P}_{\mathcal{F}} \Leftrightarrow \mathbf{AD}(c)$.

The proof has essentially two main steps. The first move, and the key one from our perspective, is to show that the set of probability functions on \mathcal{F} is the “convex hull” of the set of omniscient credence functions. That is, it’s equivalent to the set of weighted sums of such functions:

$$\mathcal{P}_{\mathcal{F}} = \mathcal{O}_{\mathcal{F}}^+ = \left\{ \sum_{w \in \mathcal{W}_{\mathcal{F}}} \lambda_w o_w(X) : \lambda_w \geq 0 \text{ and } \sum_{w \in \mathcal{W}_{\mathcal{F}}} \lambda_w = 1 \right\}$$




Figure 1: The simplest case. The dotted line is the convex hull between the two omniscient credence functions o_A (the world where A is true) and $o_{\bar{A}}$ (the world where it is false). Note that c^* is closer to any point on the line than c is. Inspired by Pettigrew (2016).

In something closer to English, what this means is that $\mathcal{P}_{\mathcal{F}}$ picks out what we can think of as the most minimal “path” between the omniscient functions. When there are only two worlds, this set will pick out a straight line segment between them; with three worlds, we get the interior of a triangle, etc. Note the important point for our purposes that the set of weighted sums ($\mathcal{O}_{\mathcal{F}}^{\pm}$) is essentially every omniscient credence function plus the various acceptable functions that mix various of these omniscient credence functions together. So change the omniscient credence functions and you change $\mathcal{O}_{\mathcal{F}}^{\pm}$. We’ll return to this point below.

The second step is to show that for any non-probabilistic credence function c , there is some probability function c^* such that for all $p \in \mathcal{P}_{\mathcal{F}}$, $\mathcal{B}(c, c^*) \leq \mathcal{B}(c, p)$. That is, there’s a probability function c^* that is the closest probability function to c . This probability function, in turn, is closer to every other probability function than c is, or, for all p , $\mathcal{B}(p, c^*) < \mathcal{B}(p, c)$. Life is easier in the simple cases (and easier still in picture, see figure 1), so think of a line segment connecting two omniscient credence functions, o_A and $o_{\bar{A}}$, (call this line “ $A\bar{A}$ ”) and a point not on the line c . Draw a second line through c perpendicular to $A\bar{A}$. The intersection of the new line with $A\bar{A}$ (call this point “ c^* ”) is both the closest point on $A\bar{A}$ to c and closer to any other point on $A\bar{A}$ than c is. The present idea is simply a generalization, and works because we’ve already proven that the set of probabilities is like a line in the sense that it’s the most minimal “path” between the omniscient functions.

Together, these two steps show is that for any non-probabilistic credence function

c , there is a probability function c^* that is closer to any probability function than c is. Since the omniscient credence functions are themselves probability functions, c^* is closer to any omniscient credence function than c is. Or: $c \notin \mathcal{P}_{\mathcal{F}} \Rightarrow \mathbf{AD}(c)$. Precisely the same results indicate that the other direction holds as well, completing the proof.

The argument from credal dominance and dominance avoidance to probabilism is usually (and I think rightly) taken to be the best argument for probabilism. As Pettigrew (2016) shows, it relies on few substantive assumptions other than that degrees of beliefs are to be judged on the basis of their accuracy. Nevertheless, beginning in the next section, I’ll argue that it fails because it relies on an assumption whose implications have been underappreciated.

2 Accuracy-dominance and logical impossibilities

In this section, I show that the accuracy-dominance argument relies on the unstated assumption that the class of worlds that accuracy is measured against does not include any worlds that are inconsistent by the lights of classical logic—that is, as I’ll usually put it, the class of worlds that accuracy is measured against is “restricted” to the classically consistent worlds.

Though the standard presentations of the accuracy dominance argument do not usually make this assumption explicit, they do implicitly assume that there are no worlds at which contradictions are true—there are no *logically* impossible worlds.⁴ This assumption is extremely important. Probability functions are by definition normalized, meaning that they assign 0 probability to contradictions. The accuracy-dominance argument relies on the fact that every omniscient credence function is a probability function and thus that every omniscient credence function assigns 0 probability to contradictions. Hence it requires that there are no worlds that violate the laws of classical logic.

We can illustrate this point by considering what happens when we introduce a world in which some contradiction is true (though exactly the same reasoning holds in the case where a tautology is false). So, suppose that there is at least one world w such that $A \wedge \neg A$ is true at this world. Then there is an omniscient credence function o_w such that $o_w(A \wedge \neg A) = 1$. Then this omniscient credence function is not normalized and thus not a probability function. Then the set of probability functions

⁴What’s explicitly assumed is a stronger condition, namely that the set of propositions form a Boolean algebra (Joyce 1998, 2009; Pettigrew 2015) or a subalgebra of such (Pettigrew 2016). Either assumption rules out logically impossible worlds and (thus) ensures that every omniscient credence function is a probability function. Note that Pettigrew (2021) drops this assumption, and with it drops probabilism as traditionally understood; see section 5 for discussion.




Figure 2: The simplest cases of accuracy dominance failure that are possible when logically impossible worlds are admitted into $\mathcal{W}_{\mathcal{F}}$. The dotted line still represents the set of probability functions, but is no longer the convex hull of the omniscient credence functions, which is now represented by the shaded area.

is not equivalent to the set of weighted sums of the omniscient credence functions. In the intuitive characterization given above, the effect is that the set of probability functions no longer picks out the most minimal “path” between omniscient credence functions (it’s not equivalent to the convex hull of such functions), and thus the geometric argument of second step of the proof fails. (In pictorial form, see figure 2.)

The above considerations show that the accuracy-dominance argument for probabilism relies on the assumption that the set of worlds includes only classically consistent worlds. This isn’t a new observation: as Williams (2012) proves, altering the logic of the worlds that accuracy is measured against alters which credence functions are picked out by the argument. If said logic is intuitionistic, then every rational credence function assigns ϕ to 1 iff ϕ can be proven constructively, and if ϕ is a constructive consequence of ψ , then there’s a familiar “no drop” principle that holds between them: an agent’s credence in ϕ should be at least as high as their credence in ψ (for details of the intuitionistic case in particular, see Weatherson 2003). Which worlds are taken to be “dominance-relevant” (Pettigrew 2021) has a substantive effect on which credence functions are rational; it’s only when the set of worlds is restricted to the classically consistent worlds that the rational credence functions are guaranteed to obey the axioms of probability. Without some motivation for restricting the dominance-relevant worlds to the classically consistent worlds, therefore, the argument does not offer any support to probabilism as traditionally understood.⁵

⁵Of course, we might simply be interested in the question of what kind of credence functions

If the accuracy-dominance argument for probabilism is to succeed, we'll need some motivation for accepting the restriction to classically consistent worlds.

The upshot is that the two premises of the accuracy-dominance argument are better represented as follows (alterations in italics):

(DA*) DOMINANCE AVOIDANCE*

An agent is rational only if they are not accuracy dominated *relative to the classically consistent worlds*.

(CD*) CREDAL DOMINANCE*

An agent's degrees of belief obey the axioms of probability if and only if they are not accuracy dominated *relative to the classically consistent worlds*.

The change from (CD) to (CD*) captures the technical point just made; the change from (DA) to (DA*) is necessary to preserve validity. The challenge facing the proponent of the accuracy-dominance argument is then to offer a motivation for accepting (DA*)—and particularly, a motivation for the restriction that distinguishes (DA*) from (DA).

3 In insufficiency of alethic considerations

In the last section, I showed that the accuracy-dominance argument for probabilism requires a commitment to (DA*): if the argument is to succeed, we need some reason to restrict the set of worlds that accuracy is measured against to the classically consistent worlds. In this section, I'll argue that alethic considerations alone are insufficient to motivate this restriction. As such, *if* there are reasons to accept (DA*) and thus the accuracy-dominance argument for probabilism, they must be based in sources of epistemic normativity other than the aim of having true beliefs. The basic reason is simple: alethic considerations—where, recall, this means considerations relating to truth modalities—alone make no distinction between a posteriori necessities and a priori ones; from the perspective of accuracy (to the truth), there's no difference

are rational given the assumption that the relevant set of worlds is the set of classically consistent worlds. I take it that this isn't what at least some proponents of the accuracy-dominance argument want out of it—the argument is supposed to provide a reason why obedience to the axioms of probability is mandated by rationality, not just an indication of what credence functions are most accurate given the assumptions.

between true beliefs about logical consequences and true beliefs about metaphysics.⁶ No matter how we try and precisify the various notions, therefore, any appeal to alethic reasons to eliminate classically impossible worlds from consideration will also motivate eliminating metaphysically impossible worlds. The result is that without further modification the accuracy-dominance argument shows that full confidence in every metaphysical necessity is a requirement of rationality—a result that I, at least, take to be unacceptable.

To illustrate the point, consider how one might argue for (DA*) on alethic grounds. If we assume that classical logic is the “one true logic” then the worlds in which classical logic is violated are not in fact possible worlds. Since there’s no sense in which classically inconsistent worlds are possible candidates for the truth, they shouldn’t be counted among the dominance-relevant worlds; there’s no situation in which these impossible worlds are the ones against which one’s beliefs are ultimately judged. In other words, the “restriction” to classically consistent worlds is in some sense no restriction at all: we’re simply measuring accuracy against all the worlds that there are. There’s no broader class of worlds that it would make sense to measure accuracy against.

The argument just sketched relies on strong assumptions about the nature of logic and rationality—we must assume both that classical logic is the one true logic and that rationality is “external and evaluative” (Williams 2018)—but it does meet the constraint that the only motivation for the restriction to classically consistent worlds is a concern with truth; there’s no appeal in this argument to concerns about evidence or justification. Nevertheless, the argument just given has a problem: it shows too much. Consider the implications of the above reasoning for an arbitrary a posteriori necessary truth ϕ . By the same argument given above, ϕ is true at every actually existing world and thus at every world that accuracy is measured against. Since ϕ is true at every dominance-relevant world, every omniscient credence function assigns a credence of 1 to ϕ . As a consequence, every member of the convex hull (that is, the set of weighted sums) of the omniscient credence functions also assigns a credence of 1 to ϕ . Then, for any credence function c that assigns a credence of less than 1 to ϕ , there is another credence function c^* that assigns a credence of 1 to ϕ and dominates c —for exactly the same reasoning as was given in the first section. The result is that assigning a credence of 1 to ϕ is a requirement of rationality, and I take it that this result is undesirable: it implies that an agent is only rational if they believe every a posteriori necessary truth.

As indicated, I intend this example to be illustrative; I don’t take it that the only possible way of justifying the restriction to classically consistent worlds on alethic

⁶Note that Pettigrew (2021, 9994) seems to anticipate this point.

grounds is precisely the argument given above. Nevertheless, this argument reveals an important and general point, namely that alethic concerns don't provide us with the resources to distinguish between a priori necessary truths—truths that are (presumably) not open to empirical evidence—and a posteriori ones that are almost by definition not propositions that agents must believe on pain of irrationality. There *cannot* be any alethic motivation for eliminating some impossibilities from consideration that doesn't tell equally in favor of eliminating all of them. The reason is simple: the distinction between an a priori necessity and a posteriori one is fundamentally an *evidential* distinction rather than an *alethic* one: both kinds of proposition are necessary; but on any view that draws this distinction, one is open to empirical evidence in a way that the other is not (compare Williams 2018, 138).⁷ The upshot is that there are no alethic grounds for the restriction found in (DA*): from a purely alethic perspective, metaphysical and logical impossibilities stand and fall together. Either both are relevant when measuring accuracy or neither is.

I think that the first of these two options is the preferable one: we should recognize that alethic considerations *alone* are insufficient to motivate interesting epistemological theses—or, even more minimally, that we cannot go directly from alethic considerations to an account of rationality without at least making a stop at the nature of evidence and justification. What we should believe is as much a matter of respecting the available evidence as it is a matter of accuracy to the truth. I'll say a bit more about this in the next section, and show how the defender of probabilism can save the accuracy-dominance argument by appealing to evidential considerations.

First, though, it's worth briefly discussing the other option, namely eliminating both metaphysical and logical impossibilities from the class of worlds against which accuracy is judged. As we've seen, the result of this restriction is that both logical and metaphysical omniscience are required by rationality. Perhaps this result does not seem like too much of a cost. After all, if we're defending probabilism, we've already accepted that rationality requires us to believe tautologies that are far beyond our ability to comprehend let alone prove; accepting the further claim that rationality requires us to believe all metaphysical necessities may seem like a relatively minimal step. It should not, however. The costs of accepting the conclusion that rationality requires full belief in all metaphysical necessities are significant and far-reaching. Consider what we might call the *enkratic constraint*: it is never permissible to believe that *A* is true while simultaneously believing that the evidence fails to support *A*—

⁷One way out of the argument just given is to reject this distinction entirely and collapse metaphysical necessity to logical necessity (perhaps plus various linguistic commitments). This seems to me like a worrying amount of metaphysics to be committed to in pursuing an epistemological thesis, but it is a potential rejoinder.

that is, it's never rational to be in a state of epistemic akrasia. Though there are some debates about this principle to be found in the literature, it's widely accepted and often taken to be a relatively fundamental or secure principle of rationality.⁸

Biting the bullet and accepting that rationality requires belief in all metaphysical necessities would require rejecting the enkratic constraint. Supposing uncontroversially that agents can lack evidence for some a posteriori metaphysical necessity, rationality will require them to nevertheless believe both that said metaphysical necessity is true and that the evidence fails to support it. Notice, importantly, that this failure of the enkratic constraint is an extremely serious one: on the view outlined, it's not just that epistemic akrasia is *permissible*; it's *required*. Further, the case in question is not one in which violations of the enkratic constraint are usually thought to arise. On the contrary, the view that full belief in metaphysical necessities is a requirement of rationality will also determine that epistemic akrasia is rationally *required* in *all* cases in which an agent believes that they fail to have definitive evidence for a metaphysically necessary proposition. We could keep on biting bullets here. But accepting that rationality requires belief in all metaphysical necessities—including in particular a posteriori necessities like “all whales are mammals”—would require a substantial rethinking of epistemology, and one that it's hard to see independent motivation for.

4 Abandoning the alethic

In the prior sections, I've argued that the accuracy-dominance argument for probabilism faces a problem: alethic considerations alone are insufficient to motivate the required restriction to classically consistent worlds. In this section, I'll turn towards a positive suggestion. Since the problem faced by the accuracy-dominance argument is that alethic considerations are insufficient to motivate a distinction between logical and metaphysical necessity—a distinction that the defender of the argument needs if they are to avoid holding that rationality requires full belief in all metaphysical necessities—the obvious solution is to appeal to non-alethic considerations. I'll illustrate the point with what seem to me to be the only plausible grounds for upholding the distinction, namely justificatory or (as I'll take to be equivalent, but will prefer because it's slightly less awkward) evidential considerations.

⁸For a survey of the literature, see Lasonen-Aarnio (2020, 600–1, note 7); for philosophers taking some form of the enkratic constraint to be fundamental or secure see, e.g., Horowitz (2014) and Titelbaum (2015, 2019). Note, however, that the considerations that motivate the enkratic constraint are not accuracy-based: it is strictly speaking open to the defender of the argument to deny it.

Here’s how this appeal might go. First, we give some story about what is required of agents when they learn evidence E . Importantly, at this juncture the story must be non-probabilistic if we want to avoid circularity—effectively, we’re only considering what worlds E makes (im)possible for the agent, not what the agent ought do in relation to all the remaining worlds. Suppose, for instance, that for the reasons outlined by defenders of logical omniscience such as Smithies (2015), we accept that the “logic of evidence” is essentially classical: E constrains what’s rational by eliminating from consideration any world that is classically incompatible with it. More precisely, we think that the right story to tell about what evidence requires is one that validates the following principle:

- (IE) (CLASSICAL) INCONSISTENCY ELIMINATION
 Learning E rules out as rationally impermissible all and only worlds classically inconsistent with E .

This principle can then be used to restrict the class of dominance-relevant worlds. Since all classical inconsistencies entail (classically, of course) a contradiction in conjunction with the empty set, this principle entails that even an agent with no evidence whatsoever must (on pain of irrationality) rule out classically inconsistent worlds. And, therefore, the dominance-relevant worlds are just the classically consistent ones: “no matter what” an agent has learned, rationality requires that their degrees of belief obey the axioms of probability. More generally, we can introduce a weakened form of (DA), as follows:

- (DA[†]) DOMINANCE AVOIDANCE[†]
 An agent is rational only if they are not accuracy dominated *relative to the rationally permissible worlds*.

In combination with (IE), this new weaker version of the dominance avoidance principle yields (DA^{*}), exactly as desired. The standard accuracy-dominance machinery then induces coherence requirements determined in the normal way: whatever is in the convex hull of the rationally permissible worlds is a rationally permissible credence function; whatever is outside it is not. Since (IE) guarantees that the dominance-relevant worlds are just the classically consistent worlds, the result is an argument in support of probabilism: (IE), (DA[†]), and (CD^{*}) jointly entail (PR).

As was true in the last section, I suspect that there are a variety of similar moves that one could make to patch the accuracy-dominance argument. My contention, however, is that if we’re to avoid the conclusion that rationality requires full belief

in all metaphysical necessities, something like the fix just given will be required. More precisely: the only way to avoid the unwanted result is to appeal to non-alethic considerations in fixing what count as dominance-relevant worlds and (plausibly) the only normative considerations that fit the bill are justificatory or evidential in the sense of relating to how the evidence constrains belief. At bottom, my reasoning here is simply that the distinction that we need to make—between a priori logical necessities and a posteriori metaphysical ones—is an evidential distinction rather than an alethic one: both kinds of necessity are true at all possible worlds, but only one is knowable without empirical evidence. Since it’s the latter property that we want to isolate, there’s no way to get around the appeal to non-alethic considerations.

Allow me to come at the issue from a slightly different direction. As I showed in §2, probabilistic credence functions dominate non-probabilistic functions because we’re only evaluating them against classically consistent worlds. This restriction could in principle be justified in a purely alethic manner—but, as we saw in the last section, any alethic justification for the restriction to classically consistent worlds seems to equally motivate a further restriction to the metaphysically possible worlds that is undesirable. So if a justification for the restriction based in purely alethic considerations is unavailable, what are the alternatives? Because the goal is to determine what is rational for *any* agent, we need something that is both (a) not alethic and (b) applies to any epistemic agent regardless of the world or context in which they find themselves. What satisfies these two desiderata are conditions that apply to all epistemic agents in virtue of being conditions on epistemic categories like belief, justification, knowledge, evidence, etc. That is: if P is a precondition of belief, then when asking the question of what it is rational to believe, we can restrict our attention to only those worlds in which P is true. The route I’ve sketched above appeals to a particular claim about the nature of evidence: to have any evidence E at all—and this includes the case where E is the empty set—is to be in a position where it is irrational to assign any credence to propositions that are (classically) inconsistent with E . This isn’t the only possible way to go, of course, but something in this vicinity is necessary: we need some doxastic or evidential principle rather than an alethic one. And if the argument for probabilism is to succeed, the relevant principle must provide us with reason to restrict the dominance-relevant worlds to those that are classically consistent (without also giving us reason to restrict them in a way that eliminates all metaphysical impossibilities from consideration).

My own view is that this project is hopeless, because it is rationally permissible to assign non-zero probability to at least some classical contradictions. So there is no true principle like (IE) that will justify the relevant restriction. Unfortunately, arguing for that conclusion would take us too far afield. Instead, I want to focus the

rest of this essay on the costs of amending the accuracy-dominance argument in a manner like the one just sketched.

5 The consequences of non-alethic considerations

In this final section, I sketch two serious consequences of adopting the amended version of the accuracy-dominance argument or something like it. The first of these is largely dialectical. Moving away from purely alethic considerations seriously changes the landscape of the argument between the probabilist who endorses the view that logical omniscience is a requirement of rationality and those epistemologists who reject this conclusion. I think this shift is already in progress: as we'll see, a number of recent papers have outlined alternatives to traditional probabilist views that drop logical omniscience; my claim is simply that whether these views are preferable is something that cannot be decided by appeal to the machinery of the accuracy-dominance argument. The second consequence is that we need to reinterpret what the accuracy-dominance argument tells us; in particular, we need a different story about what "accuracy" means in the current context and why it is something that we should value.

Beginning with the dialectical consequence. The defender of probabilism has always been committed to the principle that rationality requires logical omniscience, as this principle is a consequence of probabilism. While many see this as a high cost to pay, the probabilist in possession of a purely alethic argument for probabilism can (rightly) point out that it is supported by otherwise plausible premises about rationality and the nature of the possibility space. As we've seen, however, that argument doesn't work. Rather than allowing for the derivation of the rational necessity of logical omniscience from plausible premises, the accuracy-dominance argument requires that we assume that rationality requires logical omniscience—otherwise, we can't restrict the dominance-relevant worlds in a way that generates the right result. As such, if we can't give independent motivation for (IE) or something like it, the accuracy-dominance argument gives us no reason to endorse probabilism. In other words: absent an account of what evidence requires of agents, we can't say what view about rationality the accuracy-dominance machinery supports.

This result shouldn't be surprising. A number of recent papers—namely Dogramaci (2018), Pettigrew (2021), and Williams (2018), the former two of which are explicitly building on both the latter and Hacking (1967)—have argued against probabilism as traditionally understood on the grounds that the account of rationality found in the standard version of the accuracy-dominance argument is inappropri-

ately “externalist and evaluative” (Williams 2018).⁹ Each of these papers presents an alternative to probabilism that aims to better respect the intuitively internalist character of rationality.

The three authors go about achieving this end in slightly different ways. Williams (2018) argues for replacing the logically possible worlds found in the accuracy-dominance argument with those that are “doxastically possible”—the worlds associated with epistemic modals—and an associated move to “doxastic probabilism”: an agent is rational only if their degrees of belief are “a convex combination of the functions that pick out the truth values of propositions at each [doxastically possible] world” (Williams 2018, 134). No further constraints on admissibility are required or imposed. Pettigrew offers a conceptually distinct but similar solution: employing Hacking’s notion of “personal probability”—“a world is personally possible for a particular individual at a particular time if by this time this individual hasn’t ruled it out” (Pettigrew 2021)—rather than that of doxastic possibility, he defends “personal probabilism,” which has much the same flavor in the sense that the constraints on credence functions are structural and (at least in principle) an agent could rationally doubt any sort of logical principle. Finally, while Dogramaci’s position is also importantly distinct in a variety of technical ways, for present purposes we can think of it as “in-a-position-to-know probabilism”: an agent is rational only if their degrees of belief are coherent according to the “logic” given by all of the propositions they are in a position to know (Dogramaci 2018, 121–22).

Besides the fact that all three solutions give up on some aspects of standard probabilism—most notably, the commitment to rationality requiring logical omniscience—what’s important is that all three do so by falling back on other epistemic notions to pick out the relevant class of worlds: Williams’ epistemic “might” and “must,” Dogramaci’s “in-a-position-to-know,” and Pettigrew’s “ruling out.” These three different notions are all essentially *evidential* in the sense described in the last section: they are different ways of cashing out what the evidence requires and permits of an agent. That is: in all three cases, the grounds appealed to in picking out the dominance-relevant worlds concern what it is that evidence requires of us, not what it is that’s true. In all three cases, therefore, the proposed rethinking of rationality involves moving away from the traditional accuracy-dominance argument by introducing evidential considerations to pick out which worlds are dominance-relevant.

The upshot of the arguments outlined in this paper is to undermine the dialectical position of the probabilist relative to these alternatives. I’ve argued that

⁹Notable as well in this respect is Carr (2017, 530), who argues that the restriction to purely alethic grounds must be abandoned on the basis of a different set of cases to the ones considered here.

the probabilist has no other way to patch the accuracy-dominance argument other than appealing to evidential considerations, but that means that these evidential considerations—principles like (IE)—must be defended against alternative views about the nature of evidence like those adopted by Dogramaci (2018), Pettigrew (2021), and Williams (2018). If we’re to have a successful and persuasive accuracy-dominance argument for probabilism, we need some independent ground for thinking that rationality requires logical omniscience. Absent such ground, the argument gives us no reason to prefer probabilism as traditionally conceived to the alternatives outlined in the last paragraph.

There is an additional, less obvious, cost to adopting the proposed patch. Introducing evidential considerations into the argument—or, really, any alteration to the set of dominance-relevant worlds—requires re-thinking what the accuracy measure is actually tracking. In the purely alethic context, it does seem plausible that we’re measuring something like closeness to the truth. In any context in which we’re allowing any sort of “impossible” world, whether metaphysical or logical, this interpretation is no longer plausible: some of the credence functions that accuracy is measured against are ones that aren’t actually candidates for truth; they’re *mere* epistemic possibilities. So our accuracy measure is now tracking something other than truth and truth alone—it’s accuracy to the possibilities that haven’t been “ruled out” yet. As Carr (2017) stresses, alterations to the notion of accuracy employed in the argument raise a new and difficult question: why is this new notion of accuracy relevant to rationality? How we answer this question will depend on precisely what we think evidence requires of the agent; Williams and Pettigrew, for instance, can simply appeal to the accuracy measure tracking how well the agent is doing by their own lights, a concept with a clear connection to at least some accounts of rationality. This same option is not open to the probabilist, however, who must offer a reinterpretation that is neither purely alethic nor purely sensitive to the agent’s own evaluation of what’s possible.

Regardless of how we end up choosing to interpret the relevant of accuracy, the upshot is that *veritist* character of the accuracy-dominance argument is made promissory at best. As discussed above, the original accuracy-dominance argument is associated with a commitment to using alethic considerations alone to pick out what counts as rational. As Pettigrew has argued, this commitment fits nicely with veritism, the view that accuracy to the truth “is the only *fundamental* epistemic virtue: all other epistemic virtues derive their goodness from their ability to promote accuracy” (Pettigrew 2016, 6). The modification suggested in the last section breaks that connection; any reason that we give for restricting the set of worlds to either the by-their-own-lights possibilities or the to classically consistent worlds must be non-

alethic. The result is that the veritist has more work to do. If they are to employ the resulting argument, they must give an account of how the evidential considerations appealed to here ultimately bottom out in alethic considerations. The literature on this subject is far too vast to address here—most reliabilist accounts of justification claim to achieve this goal—but incorporating a specific account of this sort is a cost, especially since the original argument appeared to function without any premises or assumptions about the relationship between truth and evidence / justification.

Breaking the connection between the veritism and the accuracy-dominance argument has consequences for those not committed to veritism as well. More common than an explicit commitment to truth as the *sole* epistemic value is a particular way of reading the accuracy-dominance argument on which there may be other epistemic values or virtues and the accuracy-dominance argument (and other arguments like it) merely shows what accuracy to the truth requires of an agent. There may be other requirements—which may interact with the demand for accuracy in various ways—but the accuracy-dominance machinery captures at least one piece of the puzzle of rationality, namely what’s required by the aim of having true beliefs.

For the same reasons just given, however, this interpretation stands in need of additional argumentation or support. If the arguments that I’ve offered are correct, then the results of accuracy-based arguments rely on what are essentially implicit assumptions about what evidence requires of agents—for probabilism to be what alethic considerations *alone* require of an agent, we need some argument from alethic considerations to (IE) or something like it. I suspect that there cannot be an argument of this sort, but I haven’t shown that here. Regardless, the present upshot is that considerations of accuracy to truth alone are less constraining than the recent literature sometimes makes them out to be: if I’m right and some assumptions about the nature of evidence are necessary for any version of the accuracy-dominance argument, we have some reason to think that purely alethic considerations allow for a broader range of positions than may previously have been recognized.

6 Conclusion

In this paper, I’ve argued that the accuracy-dominance argument for probabilism relies on an assumption that cannot be motivated on purely alethic grounds. In particular, the argument requires restricting what Pettigrew calls the “dominance-relevant worlds” to include only those that are classically consistent. Since there are no alethic distinctions to be made between different kinds of impossibilities, any alethic justification for eliminating classical impossibilities from consideration would also require eliminating metaphysical possibilities—with the implication that rationality

requires agents not just to be logically omniscient, but also to believe all metaphysically necessary propositions as well. To avoid this unacceptable result, the only option is to fall back on evidential considerations, to build into the argument an account of what evidence requires of agents. As we've seen, there are a number of different views that are possible in this respect: evidence may require very little of agents, or it may require very much. In the former case, we're left with something weaker than probabilism; in some versions of the latter, we can still get the result that obedience to the axioms of probability is a requirement of rationality.

References

- Carr, Jennifer Rose (2017). Epistemic Utility Theory and the Aim of Belief. *Philosophy & Phenomenological Research* 95.3: 511–34.
- (2021). A Modesty Proposal. *Synthese* 198.4: 3581–601.
- Dogramaci, Sinan (2018). Solving the Problem of Logical Omniscience. *Philosophical Issues* 28: 107–28.
- Hacking, Ian (1967). Slightly More Realistic Personal Probability. *Philosophy of Science* 34.4: 311–25.
- Horowitz, Sophie (2014). Epistemic Akrasia. *Nous* 48.4: 718–44.
- Joyce, James M. (1998). A Nonpragmatic Vindication of Probabilism. *Philosophy of Science* 65.4: 575–603.
- (2009). Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief. In: *Degrees of Belief*. Ed. by Franz Huber and Christoph Schmidt-Petri: 263–97.
- Lasonen-Aarnio, Maria (2020). Enkrasia or Evidentialism? Learning to Love Mismatch. *Philosophical Studies* 177.3: 597–632.
- Pettigrew, Richard (2015). Epistemic Utility Arguments for Probabilism. In: *Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. URL: <https://plato.stanford.edu/entries/epistemic-utility/>.
- (2016). *Accuracy and the Laws of Credence*. Oxford: Oxford University Press.
- (2021). Logical Ignorance and Logical Learning. *Synthese* 198.10: 9991–10020.
- Smithies, Declan (2015). Ideal Rationality and Logical Omniscience. *Synthese* 192.9: 2769–93.
- Titelbaum, Michael (2015). Rationality's Fixed Point (or: In Defense of Right Reason). *Oxford Studies in Epistemology* 5: 253–294.
- (2019). Return to Reason. In: *Higher-Order Evidence: New Essays*. Ed. by Mattias Skipper and Asbjørn Steglich-Petersen. Oxford: Oxford University Press: 226–45.

- Weatherson, Brian (2003). From Classical to Intuitionistic Probability. *Notre Dame Journal of Formal Logic* 44.2: 111–23.
- Williams, J. Robbie (2012). Gradational Accuracy and Nonclassical Semantics. *The Review of Symbolic Logic* 5.4: 513–37.
- (2018). Rational Illogicality. *Australasian Journal of Philosophy* 96.1: 127–41.