

ACCURACY AND THE CREDENCE-BELIEF CONNECTION

Richard Pettigrew

*Department of Philosophy
University of Bristol, UK*

© 2015, Richard Pettigrew

*This work is licensed under a Creative Commons
Attribution-NonCommercial-NoDerivatives 3.0 License*

<www.philosophersimprint.org/015016/>

Probabilism is the thesis that an agent is rational only if her credences are probabilistic. This paper will be concerned with what we might call the Accuracy Dominance Argument for Probabilism (Rosenkrantz, 1981; Joyce, 1998, 2009). This argument begins with the claim that the sole source of epistemic value for a credence is its accuracy — a credence in a true proposition is more accurate, and thus of greater epistemic value, the higher it is; a credence in a false proposition is more accurate, and thus of greater epistemic value, the lower it is. The argument then proceeds to lay down properties that any numerical measure of the accuracy of credences must have. Finally, it is shown that, relative to any measure with those properties, the following holds: for any non-probabilistic credences, there are alternative credences over the same propositions that are guaranteed to be more accurate. In the language of decision theory, these alternative credences accuracy-dominate the non-probabilistic credences; they are more accurate however the world turns out to be. Thus, one can tell *a priori* that they are more accurate and thus have greater epistemic value. From this, the argument concludes, it follows that non-probabilistic credences are irrational, which is exactly what Probabilism says.

In this paper, I wish to identify and explore a lacuna in this argument. I grant that, if the only doxastic attitudes are credal attitudes, the argument succeeds. But many philosophers take this not to be the case. They say that, alongside credences, there are other doxastic attitudes, such as full beliefs, full disbeliefs, and suspensions of judgment — to wit, categorical doxastic attitudes, by contrast with credences, which are graded doxastic attitudes (Kyburg, 1961; Foley, 1992; Leitgeb, 2014; Fitelson, ms). What's more, those philosophers typically claim, these other doxastic attitudes are closely connected to credal attitudes, either as a matter of necessity — an agent has a belief in a given proposition just in case she has credences with particular properties, for instance — or normatively — for instance, if an agent has a belief in a given proposition, she ought to have credences with particular properties. Now, since categorical doxastic attitudes are also doxastic attitudes, it seems that, like credences, the sole source of epistemic value for those

attitudes is their accuracy. Thus, if we wish to measure the epistemic value of an agent's total doxastic state, we must include not only the accuracy of her credences, but also the accuracy of any full beliefs, disbeliefs, and suspensions of judgment that she has as well. However, if this is the case, there is a problem for the Accuracy Dominance Argument for Probabilism. After all, for all the argument says, there might be credences that violate Probabilism of which the following hold:

- (i) There are alternative credences in the same propositions that accuracy-dominate those credences.
- (ii) There is no *total* doxastic state — which includes categorical attitudes as well as credal attitudes — that accuracy-dominates the *total* doxastic state to which those credences belong.

After all, for all the argument says, there may be credences that are accuracy-dominated, but which are parts of total doxastic states that are not themselves accuracy-dominated by any other total doxastic state. Now, one might wonder how that could possibly happen: if one part of a total doxastic state is accuracy-dominated, surely the total state is dominated by the total state that results from replacing the dominated part by something that dominates that part and leaving everything else untouched. However, as I noted above, many philosophers think that the different parts of a total doxastic state are closely linked. It may not be possible to replace one part of a total state in a certain way without changing the rest of the total state. For instance, if there is a necessary connection between credence and belief, it may not be possible to change an agent's credences in a particular way without thereby changing her beliefs. Or, even if there is no necessary connection and it is possible to change her credences without changing her beliefs, it may be that, because there is a normative connection between the two, while the parts of the original total doxastic state — the one that includes the dominated credences — fit together in the way that rationality requires them to fit together, the parts of the new doxastic state — the one that includes the dominating credences — do not. In either case, the Accuracy Dominance Argument for Probabilism is

weakened.

The foregoing describes a possibility that the Accuracy Dominance Argument seems to leave open. Whether it in fact does leave it open is something we will answer precisely below, where we will also survey possible responses to the problem. In §1, we present the Accuracy Dominance Argument for Probabilism; in §2, we describe the lacuna in this argument that arises if we ignore the possibility of doxastic states other than credences; in the remaining sections §§3-6, we explore various ways we might try to fix the argument.

1. The Accuracy Dominance Argument for Probabilism

We represent the credal part of an agent's doxastic state by her *credence function* c . This is defined on \mathcal{F} , the set of propositions towards which the agent has a doxastic attitude. c takes each proposition X in \mathcal{F} and returns a real number $0 \leq c(X) \leq 1$ that measures her credence or degree of belief in X . Probabilism is then the following law of credences:

Probabilism An agent is rational only if her credence function is probabilistic.

The Accuracy Dominance Argument for Probabilism attempts to establish this law. According to this argument, accuracy is the sole source of epistemic value for credences. The accuracy of a credence in a true proposition increases as the credence increases; and the accuracy of a credence in a false proposition increases as the credence decreases. This allows us to compare the accuracy of two credences, but it doesn't allow us to measure that accuracy and nor does it allow us to compare two credence functions for accuracy. To do that, we need a measure of the accuracy of an individual credence at a possible world, and also a way of aggregating those individual inaccuracies to give the total inaccuracy of a credence function at a possible world. I will begin by describing the most popular way of doing this and I will run the Accuracy Dominance Argument using that; afterwards, I will show how it works for a wide variety of alternative ways of doing this.

In what follows, it will be easier to talk of the *inaccuracy* of credences and measures of that quantity rather than accuracy and measures of that. I take accuracy simply to be inaccuracy with the sign reversed. Thus, if \mathfrak{J} is a measure of inaccuracy, $-\mathfrak{J}$ is a measure of accuracy.

The most popular measure of the inaccuracy of an individual credence at a world is the so-called *quadratic scoring rule*. Let $q(i, x) := (i - x)^2$, where $i = 0$ or 1 , and $0 \leq x \leq 1$. Then $q(1, x)$ gives the quadratic scoring rule's measure of the inaccuracy of credence x in a true proposition, while $q(0, x)$ gives its measure of the inaccuracy of credence x in a false proposition. We then define the inaccuracy of a credence function at a world as the sum of the inaccuracies at that world of the individual credences it assigns — this is known as the *Brier score*. That is, if c is a credence function on \mathcal{F} and w is a possible world, let:

$$\mathfrak{B}(c, w) := \sum_{X \in \mathcal{F}} q(w(X), c(X))$$

where $w(X)$ is the indicator function of w — that is, $w(X) = 1$ if X is true at w ; $w(X) = 0$ if X is false at w . Then $\mathfrak{B}(c, w)$ is the Brier score of c at w .

Now consider an agent whom we'll call Cleo. $\mathcal{F}_{\text{Cleo}} = \{X, \bar{X}\}$. That is, Cleo has opinions only about a proposition X and its negation. And suppose her credence function is:

$$\begin{aligned} c_{\text{Cleo}}(X) &= 0.7 \\ c_{\text{Cleo}}(\bar{X}) &= 0.6 \end{aligned}$$

So Cleo violates Probabilism — c_{Cleo} is not a probability function, since the credences it assigns to X and \bar{X} do not sum to 1. Then it turns out that there are credence functions that are more accurate than c_{Cleo} regardless of whether X is true or false. Figure 1 illustrates the point.




Figure 1: We plot a credence function c defined on $\mathcal{F} = \{X, \bar{X}\}$ as the point $(c(X), c(\bar{X}))$ in the unit square. Thus, $c_{\text{Cleo}} = (0.7, 0.6)$. We also plot the indicator functions w_1 and w_2 of the two possible worlds: X is true at w_1 and false at w_2 . They are $(1, 0)$ and $(0, 1)$, respectively. Represented in this way, the Brier score of a credence function c at w_i is the square of the Euclidean distance that the credence function lies from the indicator function of w_i . The diagonal line represents the set of credence functions on \mathcal{F} that satisfy Probabilism — it is precisely those functions c for which $c(X) + c(\bar{X}) = 1$. The lower-right and upper-left blue arcs represent the credence functions that are exactly as inaccurate as c_{Cleo} at w_1 and w_2 , respectively. The credence functions that lie inside the overlap between those two arcs are those that accuracy-dominate c_{Cleo} — they are less inaccurate at w_1 and less inaccurate at w_2 . $c^* = (0.55, 0.45)$ is one of them.

Now, on its own, this doesn't show that c_{Cle0} is irrational. Some of the credence functions that dominate c_{Cle0} are themselves dominated. If all of them were, c_{Cle0} wouldn't be irrational — there is nothing irrational about being dominated by an option that is itself dominated (Pettigrew, 2013, §3). However, there are some credence functions that dominate c_{Cle0} but aren't themselves dominated. In fact, the ones that aren't themselves dominated are precisely those amongst them that satisfy Probabilism. In Figure 1, c^* is such a credence function. Thus, c_{Cle0} is irrational: accuracy is the sole source of epistemic value; and there are credence functions that have greater accuracy than c regardless of how the world turns out; and some of these credence functions aren't themselves accuracy-dominated in this way. This is one version of the accuracy dominance argument for the irrationality of c_{Cle0} : it is the version that assumes that the inaccuracy of a credence function is given by its Brier score.

The following theorem shows that the argument generalises to give an argument for the irrationality of any non-probabilistic credence function:

Theorem 1 (de Finetti)

(I) *If c is non-probabilistic, there is a probabilistic c^* such that, for all worlds w ,*

$$\mathfrak{B}(c^*, w) < \mathfrak{B}(c, w)$$

(II) *If c is probabilistic, there is no $c^* \neq c$ such that, for all worlds w ,*

$$\mathfrak{B}(c^*, w) \leq \mathfrak{B}(c, w)$$

This gives one version of the Accuracy Dominance Argument for Probabilism — it is the version that assumes that the inaccuracy of a credence function is given by its Brier score. But there is a stronger version that gains its strength by weakening that assumption. On this stronger version, all that is assumed is that the inaccuracy of a credence function is given by an inaccuracy measure generated by a continuous, strictly

proper scoring rule: the quadratic scoring rule is a continuous, strictly proper scoring rule, and it generates the Brier score, but there are many others. I will give the definition here, but I will do little to motivate the claim that inaccuracy measures ought to be generated in such a way, since that will have little to do with the arguments I give below.

A *scoring rule* is a function $\mathfrak{s} : \{0, 1\} \times [0, 1] \rightarrow [0, \infty]$. As in the case of the quadratic scoring rule, we take $\mathfrak{s}(1, x)$ to measure the inaccuracy of the credence x in a true proposition and $\mathfrak{s}(0, x)$ to measure the inaccuracy of x in a false proposition. A scoring rule is *continuous* if $\mathfrak{s}(1, x)$ and $\mathfrak{s}(0, x)$ are continuous functions of x . A scoring rule is *strictly proper* if the following holds: for all $0 \leq p \leq 1$,

$$p\mathfrak{s}(1, x) + (1 - p)\mathfrak{s}(0, x)$$

is uniquely minimised (as a function of x) at $x = p$. That is,

$$p\mathfrak{s}(1, p) + (1 - p)\mathfrak{s}(0, p) \leq p\mathfrak{s}(1, x) + (1 - p)\mathfrak{s}(0, x)$$

for all $0 \leq x \leq 1$, with equality iff $x = p$. Thus, a strictly proper scoring rule makes probabilistic credences immodest: any probabilistic credences in X and \bar{X} — that is, any p and $1 - p$, respectively, for some $0 \leq p \leq 1$ — will expect the inaccuracy of the credence in X — that is, p — to have lower inaccuracy than they will expect any other credence in X to have. Joyce (2009) argues that this is a desirable feature of an inaccuracy measure. Given a continuous, strictly proper scoring rule \mathfrak{s} , we define an inaccuracy measure $\mathfrak{I}_{\mathfrak{s}}$ for credence functions as follows

$$\mathfrak{I}_{\mathfrak{s}}(c, w) := \sum_{X \in \mathcal{F}} \mathfrak{s}(w(X), c(X))$$

We say that $\mathfrak{I}_{\mathfrak{s}}$ is *generated by \mathfrak{s}* . Here are two further examples of continuous, strictly proper scoring rules:

- *Logarithmic scoring rule*

- $l(1, x) = -\log x$
- $l(0, x) = -\log(1 - x)$

• *Spherical scoring rule*

- $\tau(1, x) = -\frac{r}{\sqrt{r^2+(1-r)^2}}$
- $\tau(0, x) = -\frac{1-r}{\sqrt{r^2+(1-r)^2}}$

We then have the following theorem:

Theorem 2 (Predd, et al., 2009) *Suppose \mathfrak{s} is a continuous, strictly proper scoring rule. And let $\mathfrak{I}_{\mathfrak{s}}$ be the inaccuracy measure it generates. Then*

- (I) *If c is non-probabilistic, there is a probabilistic c^* such that, for all worlds w ,*

$$\mathfrak{I}_{\mathfrak{s}}(c^*, w) < \mathfrak{I}_{\mathfrak{s}}(c, w)$$

- (II) *If c is probabilistic, there is no $c^* \neq c$ such that, for all worlds w ,*

$$\mathfrak{I}_{\mathfrak{s}}(c^*, w) \leq \mathfrak{I}_{\mathfrak{s}}(c, w)$$

Thus, not only is every non-probabilistic credence function accuracy-dominated by undominated alternatives when accuracy is measured by the Brier score; they are also all accuracy-dominated by undominated alternatives when accuracy is measured by an inaccuracy measure that is generated by a continuous, strictly proper scoring rule.

We are now in a position to state the strongest version of the Accuracy Dominance Argument for Probabilism:

- (1) **Veritism** Accuracy is the sole source of epistemic value for a doxastic state.
- (2) **Strict Propriety** The inaccuracy of a credence function is measured by a continuous, strictly proper scoring rule.
- (3) **Dominance** It is irrational to adopt an option that is strictly dominated by an alternative option that is not itself even weakly domi-

nated.

- (4) **Theorem 2.**

Therefore,

- (5) **Probabilism** An agent is rational only if her credence function is probabilistic.

2. A lacuna in the argument

To illustrate the lacuna in the Accuracy Dominance Argument for Probabilism, I will sketch a response that Cleo might make to the charge of irrationality that is brought against her by that argument. Recall: Cleo has credence 0.7 in X and 0.6 in \bar{X} . Her imagined response has three parts: the first is a claim about the connection between credal attitudes and categorical attitudes, such as full belief, full disbelief, and suspension of judgment; the second is a claim about how the inaccuracy of an agent's total doxastic state ought to be measured when that state consists of credences and full beliefs; the third is the observation that her total doxastic state, which includes her credence function c_{Cleo} , as well as her categorical doxastic attitudes, is not accuracy-dominated given the inaccuracy measure for total doxastic states that she has described in the second part.

2.1 The Lockean Thesis

The first part of Cleo's imagined response is the so-called Lockean Thesis (Foley, 1992). Roughly, the Lockean Thesis says this: an agent has a belief in a proposition just in case she has a sufficiently high credence in it; she has a disbelief in it just in case she has a sufficiently low credence in it; and otherwise she suspends judgment on it. To make it precise, we need to say what counts as sufficiently high and sufficiently low credence; and we need to say what modal strength the claim has. To address the first, we simply fix a threshold $\frac{1}{2} < t \leq 1$ and state the Lockean Thesis as follows, where an agent's categorical attitudes — her beliefs, disbeliefs, and suspensions — are represented by her *belief function* $b : \mathcal{F} \rightarrow \{B, D, S\}$; thus, if an agent has belief

function b , we have $b(X) = B$ iff she believes X , and so on.

Lockean Thesis with threshold t (LT[t]) If an agent has credence function c and belief function b , then:

$$\begin{aligned} c(X) > t &\Rightarrow b(X) = B \\ 1 - t < c(X) < t &\Rightarrow b(X) = S \\ c(X) < 1 - t &\Rightarrow b(X) = D \\ c(X) = t &\Rightarrow b(X) = B \text{ or } S \\ c(X) = 1 - t &\Rightarrow b(X) = D \text{ or } S \end{aligned}$$

Thus, if an agent's credence exceeds the threshold for full belief, she has a full belief; if it lies on the threshold, she has either a full belief or she suspends judgment. Similarly for full disbelief. And in all other cases, she suspends judgment. We say that a total doxastic state (b, c) consisting of belief function b and credence function c is a *Lockean_t state* if it satisfies LT[t].

Now we must consider the modal strength of the Lockean Thesis. There are, I think, three versions that it will be useful to distinguish:

- **LT[t] (Analytic)** It is analytic that LT[t].
- **LT[t] (Metaphysical)** It is metaphysically necessary that LT[t].
- **LT[t] (Normative)** It is normatively required that LT[t].

Thus, on the analytic version, beliefs are nothing over and above sufficiently high credences: to say that someone believes a proposition is simply to say that they have a sufficiently high credence in it. On the metaphysical version, by contrast, credences and categorical attitudes are distinct existences, but there is a metaphysically necessary connection between them. Finally, on the normative version, the distinct existences have no necessary connection, but there is a normative connection: to have a high credence but no belief is metaphysically possible, but it is irrational; and similarly for disbelief and suspension.

I will state Cleo's response to the charge of irrationality in terms of the metaphysical version of the Lockean Thesis — **LT[t] (Metaphysical)**

— but I will consider whether she can make a similar response by appealing to the other, weaker versions in later sections of the paper. I will conclude that the metaphysical and normative versions serve her well, while the analytic does not. Indeed, I will argue that the proponent of the Accuracy Dominance Argument for Probabilism who accepts any version of the Lockean Thesis ought to endorse the analytic version formulated above.

It is an interesting question whether any analogous problem for the Accuracy Dominance Argument for Probabilism arises for accounts of the relationship between credence and belief other than the Lockean Thesis. However, it is difficult to answer since some of the most promising such accounts apply only to probabilistic agents: that is, they only assert credal conditions on belief, disbelief, and suspension of judgment if the credences in question are probabilistic (Arló-Costa and Pedersen, 2012; Leitgeb, 2014).

2.2 *Measuring the accuracy of full beliefs*

According to the first part of Cleo's response, an agent who has credences also has categorical attitudes, such as beliefs, disbeliefs, and suspensions of judgment. Thus, her doxastic state is richer than is represented in the Accuracy Dominance Argument for Probabilism. Now, that wouldn't matter if the other parts of her doxastic state made no contribution to the overall epistemic value of her total doxastic state — the state that includes credal and categorical attitudes. But, in the second part of her response, Cleo claims that they do. As a proponent of the metaphysical version of the Lockean Thesis, Cleo maintains that categorical attitudes and credal attitudes are distinct existences. While they are intimately linked by the Lockean Thesis, they nonetheless contribute differently to the inaccuracy of any total doxastic state of which they are a part. As we saw above, the inaccuracy of a credence function at a world is measured by an inaccuracy measure that is generated by a continuous, strictly proper scoring rule. We turn now to saying how we measure the inaccuracy of an agent's belief function at a world.

As before, we begin by measuring the inaccuracy of each individual attitude; then we sum them up. Thus, we say how inaccurate is a belief in a true proposition, a suspension in a false proposition, and so on. The idea is that there is an inaccuracy that attaches to ‘getting things wrong’ — following Easwaran (ms), we denote it W and it attaches to a false belief and true disbelief — and an inaccuracy that attaches to ‘getting things right’ — again following Easwaran (ms), we denote it R and it attaches to a true belief and a false disbelief — and an inaccuracy that attaches to suspending judgment in a proposition, regardless of its truth or falsity — we denote it N . Thus, we define the function $i : \{0, 1\} \times \{B, D, S\} \rightarrow [0, 1]$ as follows — this is the analogue of a scoring rule for credences:

- $i(1, B) = i(0, D) = R$
- $i(0, B) = i(1, D) = W$
- $i(1, S) = i(0, S) = N$

Throughout most of the paper, we will make the natural assumption that $R = 0$ and $W = 1$.¹

As before, we define the inaccuracy measure for belief functions that is generated by i as follows:

$$\mathcal{I}_i(b, w) := \sum_{X \in \mathcal{F}} i(w(X), b(X))$$

We say that \mathcal{I}_i is generated by i .

Now, suppose our agent’s credence function is c and her belief function is b . And suppose s is our measure of the inaccuracy of a credence, while i is our measure of the inaccuracy of a categorical attitude. Then the inaccuracy of her total doxastic state (b, c) is defined as follows:

$$\mathcal{I}_{i,s}((b, c), w) := \mathcal{I}_i(b, w) + \mathcal{I}_s(c, w)$$

1. Since any other values can be transformed into these by a positive linear transformation, we do not lose any generality by doing this.

2.3 Cleo undominated

The final part of Cleo’s response to the Accuracy Dominance Argument brings the first and second part together. Cleo notes a particular mathematical fact: there is a Lockean threshold t , an inaccuracy measure for credence functions \mathcal{I}_s , and an inaccuracy measure for belief functions \mathcal{I}_i such that, if $(b_{\text{Cleo}}, c_{\text{Cleo}})$ is the Lockean $_t$ state that includes c_{Cleo} as a part, then there is no Lockean $_t$ state (b^*, c^*) that accuracy-dominates $(b_{\text{Cleo}}, c_{\text{Cleo}})$ when accuracy is measured by $\mathcal{I}_{i,s}$. Thus, while there are alternative credences that Cleo may adopt that are guaranteed to make the credal part of her doxastic state more accurate, by adopting them — or indeed by adopting *any* other credences — she would thereby adopt a total doxastic state that does not dominate hers; that is, she would adopt a total doxastic state that is more inaccurate than hers in at least one possible world. She submits that, while this does not ensure that her credences are rational, it does show that any irrationality does not stem from considerations of accuracy dominance.

Let’s consider Cleo’s response in a little more detail. Here is a result on which she might rely.

Theorem 3

- Let $t = 0.7$. That is, the Lockean threshold for belief is 0.7 and the threshold for disbelief is 0.3.
- Let the inaccuracy of a credence be measured by the quadratic scoring rule q ; and let the inaccuracy of a credence function be measured by the Brier score it generates $\mathfrak{B} = \mathcal{I}_q$.
- Let the inaccuracy of a categorical doxastic state be measured by i , where $R = 0$, $N = 0.3$, and $W = 1$; and let the inaccuracy of a belief function be measured by the inaccuracy measure it generates, namely, \mathcal{I}_i . Thus, the score of a maximally accurate credence — i.e. credence 1 in a truth or credence 0 in a falsehood — is the same as the score of a maximally accurate categorical state — i.e. a true belief or a false disbelief: it is 0. And similarly for a maximally inaccurate credence and a maximally inaccurate categorical

state: both score 1.

Then there is no Lockean_t state (b^*, c^*) that accuracy-dominates $(b_{\text{Cleo}}, c_{\text{Cleo}})$ when accuracy is measured by $\mathcal{J}_{i,q}$.

The theorem is proved in the Appendix (§8). The following gives an idea of how the proof goes: All of the credence functions that dominate Cleo’s credence function assign credences to X and \bar{X} that lie between $t = 0.7$ and $1 - t = 0.3$. Thus, they all belong to Lockean_{0.7} states that assign suspension of judgment to both propositions. On the other hand, Cleo’s credences belong to a Lockean_{0.7} state that includes full belief in X and suspension in \bar{X} . Thus, when X is true, Cleo’s belief function is more accurate than the belief function corresponding to any credence function that accuracy-dominates hers. What’s more, it turns out, the gain in accuracy obtained by Cleo’s categorical attitudes if X is true outweighs the loss in accuracy suffered by her credal attitudes; and so the total doxastic state is not dominated. Figure 2 illustrates this.

This, then, is Cleo’s defence against the charge of irrationality. In the remainder of the paper, I will consider objections to this response. As we will see, it proves to be quite resilient. I will begin (in §3) by considering the claim that, just as it is irrational to have a total doxastic state that is accuracy-dominated, so it is irrational to have a total doxastic state *some part of which* is accuracy-dominated. Then I will consider how Cleo’s response fares in the presence of the different versions of the Lockean Thesis formulated above (§4-5). Finally, I will ask how sensitive Cleo’s response is to the choice of threshold (§6.1) and inaccuracy measure (§6.2) that was made in Theorem 3, and to the formulation of the Lockean Thesis $\text{LT}[t]$ (§6.3).

3. Is partial dominance irrational?

We are imagining that Cleo responds to the Accuracy Dominance Argument against her by saying that, while her credences are accuracy-dominated, the total doxastic state of which they are a part is not. However, a natural objection to this response is to claim that, just as it is




Figure 2: The dashed lines represent the Lockean thresholds: the horizontal lines mark the thresholds for belief and disbelief in \bar{X} ; the vertical lines mark the thresholds for belief and disbelief in X . Note that all credence functions that dominate c_{Cleo} correspond to the belief function that assigns suspension to X and to \bar{X} .

clearly irrational to have a total doxastic state that is accuracy dominated, so it is also irrational to have a total doxastic state *some part of which is accuracy dominated*. And this is especially so in situations in which there are alternative total states no part of which is accuracy dominated. Now, that is the situation in which Cleo finds herself, since it turns out that, if (b, c) is a Lockean_t state and c is probabilistic, there is no (b^*, c^*) some part of which dominates the corresponding part of (b, c) .² So she isn't forced to be in the situation of being dominated in part; there are available options that are not. Nonetheless, she has picked an option that is. Does this make her irrational? I think not.

Consider the following situation: I must choose Option 1 or Option 2. Whichever I choose, I'll receive a gift from my friend Phil and a gift from my friend Rachel; what the gift is depends on which world we're in, w_1 or w_2 . Here are the options:

- Option 1:

	w_1	w_2
Phil	£5	£5
Rachel	£10	£10

- Option 2:

	w_1	w_2
Phil	£10	£10
Rachel	£20	£2

2. The reason is that, as we will see in §4 below, if (b, c) is Lockean_t and c is probabilistic, then c expects the inaccuracy of b , and thus the inaccuracy of (b, c) to be minimal, if the inaccuracy of b is measured by i (with $i(i, S) = N = 1 - t$) and the inaccuracy of c measured by a strictly proper scoring rule s . If an option has minimal expected inaccuracy, it cannot be accuracy dominated.

Suppose I pick Option 1. Now, part of the outcome of that option — namely, the gift bestowed by Phil — is dominated by the corresponding part of Option 2: on Option 2, Phil gives me £10 for sure; on Option 1, he gives me £5 for sure. Moreover, no part of Option 2 is dominated by the corresponding part of Option 1: neither Phil's gift nor Rachel's is guaranteed to be better on Option 1 than on Option 2. Nonetheless, it is clear that my choice of Option 1 is not thereby irrational. The upshot: being dominated in part is not sufficient for irrationality, even when the option that dominates in part is not itself dominated in part.

4. The normative version of the Lockean Thesis

The Lockean Thesis is often understood, not as having any modal force, but rather as having normative force. This is the content of **LT**[t] (**Normative**) from above. It says that credal and categorical attitudes are distinct existences that are not tied together by any modality, but nonetheless are rationally required to interact in the way stated in the Lockean Thesis. If we adopt this version of the Lockean Thesis, what becomes of Cleo's response?

We know that Cleo's total doxastic state is not dominated by any Lockean_t state, but is dominated by a non-Lockean_t state. Thus, in the presence of **LT**[t] (**Normative**), Cleo is dominated by an irrational state, but not by any rational state. Thus, our question is this: Is it sufficient for irrationality to be dominated by an irrational option?

At first, it might seem that it is. After all, if an option is dominated by an irrational option, this is surely an even more serious indictment than if it is dominated only by a rational one. If even the irrational options are guaranteed to be better than yours, you must be doing *really* badly. But this is too quick. Such an argument might go through if the irrationality of non-Lockean_t states has its source solely in their inaccuracy. If we know that one state is dominated by another that is irrational, and we know that the dominating state is irrational because it is suboptimal with respect to accuracy, then we *may* be able to infer that the original state is also suboptimal with respect to its accuracy and therefore irrational. However, if the irrationality of a non-Lockean_t

state has its source elsewhere, then the inference is not warranted.

I'll begin by expanding on this latter point. Then I will consider how $LT[t]$ (**Normative**) might follow from accuracy considerations. I will argue that, in neither case is dominance by a non-Lockean_t (and thus irrational) state sufficient for irrationality. Thus, in the presence of $LT[t]$ (**Normative**), I submit, Cleo's response succeeds.

Let us consider, then, the case in which $LT[t]$ (**Normative**) is true, but the irrationality of violating it comes not from considerations of accuracy, but from elsewhere. Let's begin by asking why, in this case, Cleo's response works. After that, we will ask from where else the normative force of the Lockean Thesis might come. Thus, we are supposing that $LT[t]$ (**Normative**) is true, but that the source of its normative force does not lie in considerations of accuracy; and we know that Cleo is dominated only by non-Lockean_t (and thus irrational) states. To see why she is not thereby irrational, consider an analogous case, namely, Philippa Foot's famous Trolley Problem (Foot, 1967; Thomson, 1976).

In the Trolley Problem, we have two options: on Option A, we change the route of the trolley, one person is killed and five are saved; on Option B, we do not change the route, five people are killed and one is saved. Which is the correct option to choose? We might reason as follows (though of course this is not the only view we might take of the Trolley Problem). Taking into account only considerations of utility, A is the correct option, since it dominates B — five lives give rise to more utility than one. But when we take into account other considerations, such as the impermissibility of intentional killing, Option A is ruled out from the start; it is impermissible before the utility considerations come into play. Thus, although B is a dominated action, it is dominated only by an alternative option that is impermissible. And this isn't sufficient to rule it out as a permissible action. The upshot of this is that the badness of being dominated depends on the nature of the dominating option. If that dominating option is impermissible for other reasons — reasons not captured by the utility function — then the option it dominates is not necessarily irrational.

Now, the utility of a doxastic state is its accuracy. Thus, an analo-

gous point applies: if one doxastic state is dominated by another that is impermissible for other reasons — reasons not encoded in the inaccuracy measure — then the state it dominates is not necessarily irrational. In particular, it is no indictment of Cleo's rationality that her state is dominated by a non-Lockean_t alternative, if the irrationality of that alternative stems from something other than its inaccuracy.³

Now, what other source of irrationality might there be? It might be, for instance, that there are basic, brute norms of doxastic consistency. Thus, it might be a brute normative fact — that is, one not justified by any more basic normative facts, and in particular not justified by appeal to accuracy considerations — that it is irrational to have a doxastic state that does not present a consistent set of attitudes towards the world. This is analogous to the Trolley Problem case, where there is a brute normative fact that acts that involve intentional killing are impermissible — this fact, we might suppose, is not justified by more basic normative facts, in particular, not by facts about utility. If that's the case, it's plausible that an agent with credence 0.9 in proposition X along with a full disbelief in X is failing to present a consistent set of attitudes towards the world: they are akin to an agent with full belief X and full disbelief in X .

Such a move would support Cleo's response by rendering the states that dominate hers irrational and thus impermissible, and thereby rendering the fact of their dominance irrelevant to her rationality. Nonetheless, I would advise against the move. After all, a key premise of the Accuracy Dominance Argument for Probabilism is Veritism, the

3. In their "evidentialist" worry about the Accuracy Dominance Argument for Probabilism, Easwaran and Fitelson (2012) claim that evidence imposes normative constraints on the credences of any agent that has that evidence, and that violating these constraints renders an agent's credence function impermissible in a way that does not depend on accuracy considerations. They then argue that there are non-probabilistic credence functions that satisfy certain evidential constraints and are dominated only by credence functions that violate those evidential constraints. They conclude that the accuracy dominance argument fails to show that these credence functions are irrational for agents with the relevant evidence since they are dominated only by impermissible credence functions. See (Pettigrew, 2013) for a response.

view that the sole source of epistemic value for doxastic states is their accuracy. The account of the irrationality of non-Lockean_t states just presented suggests that there is some extra source of value so that value accrues to a doxastic state in virtue of it being consistent in a certain way. If that is genuinely *extra* value — that is, if it is not value that already accrues to the state in virtue of its accuracy — then Veritism must be false. Thus, if Cleo wishes to appeal to **LT**[*t*] (**Normative**) to defend herself against the Accuracy Dominance Argument, she will need to justify it by appeal to the virtue of accuracy alone.

Now, as Carl Hempel observed (Hempel, 1962, §12), and as Kenny Easwaran has developed in detail (Easwaran, ms), there is an accuracy-based justification of **LT**[*t*] (**Normative**). I will describe this now and ask whether Cleo might appeal to it in her response to the Accuracy Dominance Argument.

Suppose an agent has a probabilistic credence function *c*. And suppose she is deciding which categorical doxastic state she should adopt. The natural suggestion is that she ought to choose a categorical state that minimises expected inaccuracy by the lights of her probabilistic credence function and relative to the inaccuracy measure for categorical states that was described above. That is, she ought to adopt a belief function *b* such that, for all belief functions *b'*, we have

$$\sum_w c(w)i(b, w) \leq \sum_w c(w)i(b', w)$$

As Hempel and Easwaran have shown, if we assume that $R = 0$ and $W = 1$ and N is closer to R than to W , then *b* minimizes expected inaccuracy by the lights of *c* iff (b, c) is a Lockean_{1-N} state.⁴ Thus, at least for agents with probabilistic credence function, **LT**[$1 - N$] (**Normative**) follows from considerations of accuracy alone. If one makes the further assumption that, for agents with the same credence in an

4. More generally, if $N = 0$ and $R < W$, then *b* minimises expected inaccuracy by the lights of *c* iff (b, c) is a Lockean _{$\frac{W}{R+W}$} state.

individual proposition, the same categorical doxastic attitudes towards that proposition are permissible, we then obtain **LT**[$1 - N$] (**Normative**) for all agents, probabilistic or not. Let us grant this. Can Cleo then use it in her response to the Accuracy Dominance Argument that charges her with irrationality? That is, can she defend her dominated state by pointing out that it is dominated only by irrational states, where their irrationality lies in their violation of **LT**[$1 - N$] (**Normative**)?

At first, and for similar reasons to those entertained above, it might seem not. For instance, we know that Cleo's state $(b_{\text{Cleo}}, c_{\text{Cleo}})$ is dominated by (b_{Cleo}, c^*) . We also know that (b_{Cleo}, c^*) violates **LT**[$1 - N$] (**Normative**) when $N = 0.3$. Thus, we know that there is a belief function $b^* \neq b_{\text{Cleo}}$ such that c^* expects b^* to be more accurate than c^* expects b_{Cleo} to be. Thus, we might think: Cleo's total state $(b_{\text{Cleo}}, c_{\text{Cleo}})$ is dominated by (b_{Cleo}, c^*) ; and, from the point of view of (b_{Cleo}, c^*) , the total state (b^*, c^*) is an optimal doxastic state. Thus, for dominance reasons, we ought to prefer (b_{Cleo}, c^*) to $(b_{\text{Cleo}}, c_{\text{Cleo}})$; and then for expected accuracy reasons, we ought to prefer (b^*, c^*) to (b_{Cleo}, c^*) . Thus, by the transitivity of preference, we ought to prefer (b^*, c^*) to $(b_{\text{Cleo}}, c_{\text{Cleo}})$. Thus, Cleo is irrational, since there is an available option that ought to be preferred to the one she has adopted.

However, there is an illegitimate move in this argument. We know that there is b^* such that c^* expects b^* to be more accurate than it expects b_{Cleo} to be. And we moved from this to the claim that, from the point of view of (b_{Cleo}, c^*) , the total state (b^*, c^*) is better than (b_{Cleo}, c^*) itself. But that move is not warranted. It is true that, from the point of view of *part of* (b_{Cleo}, c^*) — namely, the credal part c^* — (b^*, c^*) is better than (b_{Cleo}, c^*) . But it does not follow from this that the same is true from the point of view of *the total state* (b_{Cleo}, c^*) . After all, if we instead ask which state is optimal from the point of view of b_{Cleo} — that is, the other part of the total state — we would most likely receive as an answer a state that includes b_{Cleo} as its belief function. And if we ask what is better from the point of view of the whole state (b_{Cleo}, c^*) , it isn't at all clear what the verdict would be.

What we have just encountered is an instance of a more general

problem that arises when one seeks guidance from a doxastic state that doesn't present a consistent attitude to the world: different parts of such a state often give rise to different and incompatible preference orderings. This is what Dutch Book arguments are often taken to dramatise (cf. the "divided mind" interpretation in [Armendt, 1993] that is criticised by Vineberg [2001]). In the Dutch Book argument, a credence in $A \vee B$ determines one preference ordering over bets on that proposition, while credences in A and in B together determine another preference ordering on essentially the same bets. Only if the credence in the disjunction is the sum of the credences in the mutually exclusive disjuncts do we obtain compatible preference orderings. This is one sense in which a non-probabilistic credence function constitutes an inconsistent attitude towards the world. Here is another similar way. Suppose c is a credence function and we have two partitions $\{w_1, w_2, w_3, w_4\}$ and $\{w_1 \vee w_2, w_3 \vee w_4\}$, where $c(w_1) + \dots + c(w_4) = 1$ and $c(w_1 \vee w_2) + c(w_3 \vee w_4) = 1$. If we wish to use this credence function to choose between two options each of which has the same utility at worlds w_1 and w_2 and at worlds w_3 and w_4 , then we can use credences over either partition to calculate the expected utility. But, unless c is probabilistic and $c(A \vee B) = c(A) + c(B)$ and $c(C \vee D) = c(C) + c(D)$, then these two expected utility calculations might give rise to different preference orderings. The same happens in the case in hand here. We cannot infer from the fact that (b^*, c^*) is better than some other state by the lights of c^* that it is better than that other state by the lights of (b_{Cleo}, c^*) . While we may infer — as the Hempel-Easwaran argument for **LT**[1 – N] (**Normative**) does in fact infer — from the fact that c expects b^* to be more accurate than it expects b to be that (b, c) is an irrational total state, we cannot infer anything about what the state (b, c) prefers, and in particular we cannot infer whether or not it prefers (b^*, c) to itself.

Thus, again, it seems Cleo's response is saved. While there are total states that dominate hers, they are all irrational. And this, we have argued, makes their existence irrelevant to Cleo's rationality.

5. The analytic version of the Lockean Thesis

So far, we have considered Cleo's response to the Accuracy Dominance Argument in the presence of two versions of the Lockean Thesis: the metaphysical and the normative versions. We have seen that Cleo's response works in the presence of both. In each of these versions of the thesis, categorical doxastic states and credal doxastic states are distinct existences. In this penultimate section, we consider the remaining version of the Lockean Thesis, namely, **LT**[t] (**Analytic**). On this version, doxastic states and credal states are not distinct existences: indeed, each doxastic state is merely a species of credal state. That is, when we say that an agent has a full belief in X , we say no more than that she has a sufficiently high credence in X . Thus, 'Cleo believes X ' is akin to 'Cleo is tall': in the latter case, the proposition says simply that Cleo's height is in a given range; likewise, in the former case, the proposition says simply that Cleo's credence in X is in a given range.⁵ Now, if this is the case — if full belief and disbelief are *reduced* to sufficiently high and low credence, respectively — it cannot be that having a true full belief lends extra accuracy to one's total doxastic state over and above the accuracy obtained from having the high credence that underlies that full belief ascription. Thus, in the presence of **LT**[t] (**Analytic**), the correct inaccuracy measure for a total doxastic state (b, c) is:

$$\mathcal{I}_s((b, c), w) := \mathcal{I}_s(c, w)$$

That is, the inaccuracy of the total state is just the inaccuracy of the credal part — the belief function adds nothing more, since it is nothing

5. Also, if we consider again how we formulated the Lockean Thesis, we see another analogy between the two propositions: the predicates in both cases — 'believes' and 'is tall' — admit borderline cases. In the latter case, if Cleo is 5'9", it is undetermined whether or not she is tall. In the former case, if Cleo has credence t in X , it is undetermined whether Cleo believes X or suspends judgment; and similarly, if she has credence $1 - t$ in X , it is undetermined whether Cleo disbelieves X or suspends judgment.

more than a coarse-grained summary of the credence function. In this case, then, Cleo's response fails: after all, her total state $(b_{\text{Cleo}}, c_{\text{Cleo}})$ is accuracy dominated by the Lockean_{*t*} state (b^*, c^*) when inaccuracy is measured in the way just described. This, I submit, is how the proponent of the Accuracy Dominance Argument for Probabilism should respond to Cleo's defence.

6. The robustness of Cleo's response

We have just seen that the proponent of the Accuracy Dominance Argument for Probabilism can avoid Cleo's response by adopting **LT**[*t*] (**Analytic**). But perhaps they need not do so. Perhaps Cleo's response depends crucially on some detail that might be questioned. In this section, we ask how robust Cleo's response is. Does it depend crucially on the particular threshold that was chosen? Or on the scoring rule used to measure the inaccuracy of the credences? Or on the inaccuracies assigned to getting it right, getting it wrong, and suspending? Or on the formulation of the non-modal part of the Lockean Thesis, namely, **LT**[*t*]?

If that were the case, then one might restore the Accuracy Dominance Argument by ruling out the Brier score as a legitimate measure of the inaccuracy of credence functions; or by showing by example that the Lockean threshold $t = 0.7$ is too low or too high; or by arguing that suspensions should be deemed more inaccurate than Cleo deems them, or less; or one might feel that belief requires credence *strictly above* the threshold by contrast with the formulation of the Lockean Thesis presented above, which permits belief at the threshold. And indeed concerns have been raised, in a different context, about the legitimacy of the Brier score (Levinstein, 2012). And it is probably true that $t = 0.7$ is too low a threshold for belief — when I roll a die, is it rational to have a full belief that it won't land 1? Moreover, one might well think that suspensions are treated too leniently by Cleo's inaccuracy measure. We will consider each of these concerns in turn.

6.1 Sensitivity to choice of threshold

The first thing to say is that, if the Lockean threshold increases, Cleo's total doxastic state will become non-Lockean relative to the new threshold. Moreover, if she retains the credal part of that state — namely, c_{Cleo} — and adopts the belief function that is required by the Lockean Thesis with the new threshold, her total state will be dominated, regardless of whether the inaccuracy of suspension changes. For, in this case, her new belief function will assign suspensions of judgment to X and to \bar{X} . And these are the same categorical attitudes to which the credence functions that dominate her credence function give rise via the Lockean Thesis. Thus, the Lockean states to which those dominated credence functions belong will dominate her total state.

However, such a move will only buy the Accuracy Dominance Argument a little time, as the following theorem shows:

Theorem 4 *For all Lockean thresholds t , there is a Lockean_{*t*} state (b, c) such that:*

- (i) *c is non-probabilistic;*
- (ii) *There is no Lockean_{*t*} state (b^*, c^*) that accuracy-dominates (b, c) when accuracy is measured by $\mathcal{J}_{i,q}$, where i is given by $R = 0, N = 1 - t, W = 1$.*

The proof is given in the Appendix (§8). Thus, while it is true that, if the Lockean threshold is shifted upwards, Cleo will become dominated, this doesn't save the Accuracy Dominance Argument for Probabilism, for there will be other non-probabilistic credence functions that belong to Lockean total doxastic states that are not dominated.

6.2 Sensitivity to choice of inaccuracy measure

Our second theorem addresses those who are concerned not with the threshold that Cleo has chosen, but with the scoring rule. To state this theorem, we need some definitions.

- We say that a scoring rule is *normalised* if $s(0,0) = s(1,1) = 0$ and $s(0,1) = s(1,0) = 1$.

- We say that a scoring rule generates a *connected* inaccuracy measure \mathfrak{I}_s if, for each credence function, the set of credence functions that dominates it is either empty or connected (in the topological sense).

The spherical and quadratic scoring rules are continuous, normalised, connected, and strictly proper; the logarithmic scoring rule is not.

Theorem 5 *Suppose s is a continuous, normalised strictly proper scoring rule that generates a connected inaccuracy measure \mathfrak{I}_s . Then there is $0.5 < r \leq 1$ such that, for all Lockean thresholds $t \geq r$, there is a Lockean $_t$ state (b, c) such that*

- c is non-probabilistic;*
- There is no Lockean $_t$ state (b^*, c^*) that accuracy-dominates (b, c) when accuracy is measured by $\mathfrak{I}_{i,s}$, where i is given by $R = 0$, $N = 1 - t$, and $W = 1$.*

The proof is given in the Appendix (§8). Thus, Cleo's strategy does not depend crucially on measuring the inaccuracy of credences using the quadratic scoring rule. For many ways of measuring the inaccuracy of credences, there are thresholds t and Lockean $_t$ states containing non-probabilistic credence functions that are undominated.

6.3 A stricter Lockean Thesis

So we have seen that Cleo's defence is not very sensitive either to the Lockean threshold we choose, nor to the inaccuracy measure we use. However, as it stands, it does seem to be sensitive to the formulation of the Lockean Thesis. In this section, we explore this sensitivity.

According to the formulation we gave above — that is, $LT[t]$ — an agent has a full belief if her credence is above the upper threshold t and a full disbelief if her credence is below the lower threshold $1 - t$; and she suspends judgment if she has a credence that lies strictly between them. But the categorical state she is in if she lies on either threshold is undetermined: if she lies on the upper threshold, then she either suspends or has a full belief; if she lies on the lower threshold, then

she either suspends or has a full disbelief. But what happens if we remove this indeterminacy at the thresholds?

There are two ways to do this: on the first, non-strict version, a credence on the upper threshold gives a full belief, and a credence on the lower threshold gives a full disbelief; on the second, strict version, a credence on either threshold results in suspension of judgment.

It's clear that Cleo's response will still work in the presence of the non-strict version. After all, Cleo's total doxastic state $(b_{\text{Cleo}}, c_{\text{Cleo}})$ satisfies that version; and the set of total states permitted by that version is a proper subset of the set of total states permitted by the original version of the Lockean Thesis. So, since Cleo is undominated in the presence of the original Lockean Thesis, she is undominated in the presence of the non-strict one.

Now, it is also the case that the set of states permitted by the strict version of the Lockean Thesis is a proper subset of the set of states permitted by the original version. But, in this case, Cleo's own state is not amongst them. That is, Cleo herself does not satisfy the strict Lockean Thesis (for the threshold $t = 0.7$) — after all, she has a full belief in X , but her credence in X lies on the threshold $t = 0.7$. So, while she is not dominated by a strictly Lockean $_t$ state, this has little force against the charge of irrationality since her own state is not itself strictly Lockean $_t$.

Suppose, however, that we consider a state $(b_{\text{Theo}}, c_{\text{Theo}})$ where c_{Theo} is not probabilistic, but the total state is strictly Lockean $_t$. Then it will turn out that Cleo's strategy for responding to the charge of irrationality that $(b_{\text{Theo}}, c_{\text{Theo}})$ faces from the Accuracy Dominance Argument will not work as it stands. The problem is that any such pair will be dominated. The crucial feature of Cleo's state in the presence of the non-strict Lockean Thesis and the original version is that all of the credence functions that accuracy dominate her credence function — that is, c_{Cleo} — belong to Lockean $_t$ states whose belief function is not b_{Cleo} . Thus, if Cleo were to move to one of those credence functions, she'd be forced to abandon her belief function, and at some worlds that loses her accuracy that is not compensated for by the accuracy

she gains by moving to the dominating credence function. This is not the case for a state $(b_{\text{Theo}}, c_{\text{Theo}})$ that satisfies the strict Lockean Thesis. At least some of the credence functions that dominate c_{Theo} belong to Lockean_{*t*} states that include b_{Theo} . Let c^* be such a credence function. So c^* accuracy dominates c_{Theo} when accuracy is measured by \mathcal{I}_s ; and (b_{Theo}, c^*) is strictly Lockean_{*t*}. Thus, (b_{Theo}, c^*) accuracy dominates $(b_{\text{Theo}}, c_{\text{Theo}})$ when accuracy is measured by $\mathcal{I}_{i,s}$, however *i* is defined. Figure 3 illustrates the point.

Is this, then, a move that the proponent of the Accuracy Dominance Argument for Probabilism might make? Might she save the argument by moving to the normative or metaphysical version of the *strict* Lockean Thesis, rather than adopting the analytic version of the *original* Lockean Thesis?

This move would certainly allow her to retain the claim that credal and categorical doxastic states are distinct existences. However, there are two problems with this move: First, the normative version of the strict Lockean Thesis is uncomfortable, at least for someone who subscribes to Veritism, as the proponent of the Accuracy Dominance Argument does. After all, as we saw in §4, it is the weaker original version for the Lockean Thesis that is justified by considerations of expected accuracy. If an agent has a credence t in X and $1 - t$ in \bar{X} , and if her inaccuracy measure for categorical doxastic states is i (where $i(i, S) = N = 1 - t$), then her expected inaccuracy for suspending judgment on X and her expected inaccuracy for having a full belief in X are equal. But the strict version of the Lockean Thesis rules out having a full belief in X as irrational, and demands that the agent suspends judgment on X . Thus, any normative version of the strict Lockean Thesis must appeal to some source of irrationality beyond accuracy considerations to explain the normative difference between full belief in X and suspension in X when they have equal expected inaccuracy. And this is what Veritism is intended to rule out.

The other problem with this move is that it only avoids the original version of Cleo’s response. Cleo noted that her Lockean_{*t*} state isn’t dominated by any other Lockean_{*t*} state. We have now seen that any




Figure 3: Suppose $(b_{\text{Theo}}, c_{\text{Theo}})$ is a strictly Lockean_{*t*} state, where c_{Theo} is as shown. Thus, $b_{\text{Theo}}(X) = B$ and $b_{\text{Theo}}(\bar{X}) = S$. Now, c_{Theo} is dominated by c^* . Therefore, (b_{Theo}, c^*) dominates $(b_{\text{Theo}}, c_{\text{Theo}})$. And it is easy to see that (b_{Theo}, c^*) is a strictly Lockean_{*t*} state as well. But notice that c^* is also dominated by further credence functions that give rise to b_{Theo} via the strict Lockean Thesis with threshold t . Indeed there is no strictly Lockean_{*t*} state that dominates $(b_{\text{Theo}}, c_{\text{Theo}})$ but is not itself dominated. In fact, this is true of any strictly Lockean_{*t*} state whose credal component is not probabilistic: it is accuracy dominated by further strictly Lockean_{*t*} states, but not by any strictly Lockean_{*t*} state that is itself undominated.

strictly Lockean_t state is dominated by another strictly Lockean_t state. But, as we saw briefly above, being dominated is not sufficient for irrationality. What is sufficient for rationality is being dominated *by an option that is not itself dominated*. And it is easy to see that, while any strictly Lockean_t state — such as $(b_{\text{Theo}}, c_{\text{Theo}})$ — is accuracy dominated, every strictly Lockean_t state that dominates it is itself accuracy dominated by another strictly Lockean_t state. Thus, it seems that those states too are not rendered irrational by the Accuracy Dominance Argument for Probabilism.

7. Conclusion

In sum: We have seen that, if categorical doxastic states — such as full belief, full disbelief, and suspension of judgment — are distinct existences, and if they are connected by the metaphysical or normative versions of the Lockean Thesis, then the Accuracy Dominance Argument for Probabilism fails: there are total doxastic states that include non-probabilistic credence functions that are not accuracy dominated when their total inaccuracy is considered. On the other hand, if they are not distinct existences and are instead connected by the analytic version of the Lockean Thesis, then the Accuracy Dominance Argument for Probabilism succeeds.

In this paper, we have considered only one other type of doxastic attitude, namely, the categorical attitudes. But there are others. For instance, comparative attitudes are captured by an ordering of propositions, *X is at most as likely as Y*, which we write $X \preceq Y$. And there is a natural analogue of the Lockean Thesis for these attitudes: $c(X) \leq c(Y)$ iff $X \preceq Y$. If these are distinct existences, they might contribute differently to the inaccuracy of the total state of which they are a part (cf. [Fitelson, ms]). And then similar issues might arise as arose in this paper. I leave open the question of how accuracy considerations for credences interact with accuracy considerations for other doxastic states.⁶

6. I am very grateful to the following people for very helpful discussion of earlier versions of this paper: Branden Fitelson, Carlotta Pavese, Kenny Easwaran,

8. Appendix: Proofs of Theorems 3, 4, and 5

In this appendix, we prove Theorems 3, 4, and 5. First, we pick our inaccuracy measures:

- Let \mathfrak{s} be a continuous, normalised strictly proper scoring rule that generates a connected inaccuracy measure $\mathfrak{J}_{\mathfrak{s}}$.
- Let i be defined as follows, where $t > 0.5$:
 - $i(1, B) = i(0, D) = R = 0$
 - $i(0, B) = i(1, D) = W = 1$
 - $i(1, S) = i(0, S) = N = 1 - t$

So

$$\mathfrak{J}_i(b, w) := \sum_{X \in \mathcal{F}} i(w(X), b(X))$$

Then the inaccuracy of a total doxastic state (b, c) is as follows:

$$\mathfrak{J}_{i, \mathfrak{s}}((b, c), w) := \mathfrak{J}_i(b, w) + \mathfrak{J}_{\mathfrak{s}}(c, w)$$

Now, let c be a credence function on $\mathcal{F} = \{X, \bar{X}\}$ such that $c(X) = t$ and $c(\bar{X}) = 1 - t + \varepsilon$ (where $0 < \varepsilon < 2t - 1$). We write this, using the vector notation, as $c = (t, 1 - t + \varepsilon)$. And let b be the belief function on \mathcal{F} such that $b(X) = B$ and $b(\bar{X}) = S$. Thus, (b, c) is Lockean_t. Moreover,

- $\mathfrak{J}_{i, \mathfrak{s}}((b, c), w_1) = \mathfrak{J}_i(b, w_1) + \mathfrak{J}_{\mathfrak{s}}(c, w_1) = \mathfrak{J}_{\mathfrak{s}}(c, w_1) + N = \mathfrak{s}(1, t) + \mathfrak{s}(0, 1 - t + \varepsilon) + 1 - t$
- $\mathfrak{J}_{i, \mathfrak{s}}((b, c), w_2) = \mathfrak{J}_i(b, w_2) + \mathfrak{J}_{\mathfrak{s}}(c, w_2) = \mathfrak{J}_{\mathfrak{s}}(c, w_2) + 1 + N = \mathfrak{s}(0, t) + \mathfrak{s}(1, 1 - t + \varepsilon) + 2 - t$

Jason Konek, an anonymous referee for this journal, and everyone involved in the symposium on Branden Fitelson's *Coherence* manuscript at Duke University in May 2014. I was supported by an ERC Starting Researcher Grant 'Epistemic Utility Theory: Foundations and Applications' during his work on this paper.




Figure 4: DB marks the set of credence functions that belong to Lockean_t states whose belief functions assign disbelief to X and belief to \bar{X} . And similarly for others.

Our next step is to divide the Lockean_t states that consist of belief and credence functions defined on \mathcal{F} into nine sets, corresponding to the nine different combinations of belief, disbelief, and suspension that an agent might have towards X and \bar{X} . Thus, we write BB for the set of Lockean_t states in which the agent believes both; we write DS for the set of Lockean_t states in which the agent disbelieves X and suspends on \bar{X} ; and so on. These are marked in Figure 4. Now we list certain facts about the relationship between the inaccuracies of (b, c) at the worlds w_1 and w_2 and the inaccuracies of a Lockean_t state that falls in one of these sets.

1. $(b^*, c^*) \in DB$:

- (a) $\mathcal{J}_{i,s}((b^*, c^*), w_1) = \mathcal{J}_s(c^*, w_1) + 2$
(This is because b^* gets both categorical doxastic states wrong at w_1 .)
- (b) $\mathcal{J}_s(c^*, w_1) > \mathcal{J}_s(c, w_1)$.
(This is because $s(1, x)$ is a decreasing function of x and $s(0, x)$ is an increasing function of x .)

2. $(b^*, c^*) \in DS$:

- (a) $\mathcal{J}_{i,s}((b^*, c^*), w_1) = \mathcal{J}_s(c^*, w_1) + 1 + N = \mathcal{J}_s(c^*, w_1) + 1 + 1 - t$
(This is because b^* gets one categorical doxastic state wrong at w_1 and suspends on the other.)
- (b) $\mathcal{J}_s(c^*, w_1) \geq s(1, 1-t) + s(0, 1-t)$
(This is because the credence function in DS that has lowest inaccuracy at w_1 is $c^* = (1-t, 1-t)$.)

3. $(b^*, c^*) \in SB$:

- (a) $\mathcal{J}_{i,s}((b^*, c^*), w_1) = \mathcal{J}_s(c^*, w_1) + 1 + 1 - t$
(This is because b^* gets one categorical doxastic state wrong at w_1 and suspends on the other.)
- (b) $\mathcal{J}_s(c^*, w_1) \geq s(1, t) + s(0, t)$
(This is because the credence function in SB that has lowest inaccuracy at w_1 is $c^* = (t, t)$.)

4. $(b^*, c^*) \in DD$:

- (a) $\mathcal{J}_{i,s}((b^*, c^*), w_1) = \mathcal{J}_s(c^*, w_1) + 1$
(This is because b^* gets one categorical doxastic state wrong at w_1 and the other right.)
- (b) $\mathcal{J}_s(c^*, w_1) \geq s(1, 1-t) + s(0, 0) = s(1, 1-t)$
(This is because the credence function in DD that has lowest inaccuracy at w_1 is $c^* = (1-t, 0)$.)

5. $(b^*, c^*) \in BB$:

- (a) $\mathcal{J}_{i,s}((b^*, c^*), w_1) = \mathcal{J}_s(c^*, w_1) + 1$
(This is because b^* gets one categorical doxastic state wrong at

w_1 and the other right.)

(b) $\mathfrak{J}_s(c^*, w_1) \geq s(0, t) + s(1, 1) = s(0, t)$

(This is because the credence function in BB that has lowest inaccuracy at w_1 is $c^* = (1, t)$.)

6. $(b^*, c^*) \in SS$:

(a) $\mathfrak{J}_{i,s}((b^*, c^*), w_1) = \mathfrak{J}_s(c^*, w_1) + 1 - t + 1 - t$

(This is because b^* suspends on both propositions.)

(b) $\mathfrak{J}_s(c^*, w_1) \geq s(1, t) + s(0, 1 - t)$

(This is because the credence function in SS that has lowest inaccuracy at w_1 is $c^* = (t, 1 - t)$.)

7. $(b^*, c^*) \in BS$:

(a) $\mathfrak{J}_{i,s}((b^*, c^*), w_1) = \mathfrak{J}_s(c^*, w_1) + 1 - t$

(b) $\mathfrak{J}_{i,s}((b^*, c^*), w_2) = \mathfrak{J}_s(c^*, w_2) + 1 + 1 - t$

8. $(b^*, c^*) \in SD$:

(a) $\mathfrak{J}_{i,s}((b^*, c^*), w_1) = \mathfrak{J}_s(c^*, w_1) + 1 - t$

(b) $\mathfrak{J}_{i,s}((b^*, c^*), w_2) = \mathfrak{J}_s(c^*, w_2) + 1 + 1 - t$

9. $(b^*, c^*) \in BD$:

(a) $\mathfrak{J}_{i,s}((b^*, c^*), w_2) = \mathfrak{J}_s(c^*, w_2) + 2$

(This is because b^* gets both categorical doxastic states wrong at w_2 .)

(b) $\mathfrak{J}_s(c^*, w_2) > \mathfrak{J}_s(c, w_2)$.

(This is because $s(1, x)$ is a decreasing function of x and $s(0, x)$ is an increasing function of x .)

Now suppose that we pick threshold $t > 0.5$ such that the following holds:

(A) $s(1, t) + s(0, 1 - t) - s(1, 1 - t) < t$

(B) $s(1, t) + s(0, 1 - t) - s(0, t) < t$

This is always possible, since $s(0, t)$ is a strictly increasing function of t and $s(1, t)$ is a strictly decreasing function of t .

Then let $b = (B, S)$ and $c = (t, 1 - t + \varepsilon)$. We wish to show the following: there is $\varepsilon > 0$ such that

- If $(b^*, c^*) \in DB, SB, DS, DD, BB, SS$, then

$$\mathfrak{J}_{i,s}((b, c), w_1) < \mathfrak{J}_{i,s}((b^*, c^*), w_1)$$

So (b^*, c^*) does not dominate (b, c) .

- If $(b^*, c^*) \in BS, SD$, then

$$\mathfrak{J}_{i,s}((b, c), w_1) < \mathfrak{J}_{i,s}((b^*, c^*), w_1)$$

or

$$\mathfrak{J}_{i,s}((b, c), w_2) < \mathfrak{J}_{i,s}((b^*, c^*), w_2)$$

So (b^*, c^*) does not dominate (b, c) .

- If $(b^*, c^*) \in BD$, then

$$\mathfrak{J}_{i,s}((b, c), w_2) < \mathfrak{J}_{i,s}((b^*, c^*), w_2)$$

So (b^*, c^*) does not dominate (b, c) .

To do this, we show that these inequalities hold for $\varepsilon = 0$. Then, noting that s is continuous, we note that there must therefore be some $\varepsilon > 0$ for which they continue to hold. Thus, note that, for $\varepsilon = 0$, and in the presence of (A) and (B), we have:

- If $(b^*, c^*) \in DB$, then, by (1),

$$\mathfrak{J}_{i,s}((b, c), w_1) = \mathfrak{J}_s(c, w_1) + 1 - t < \mathfrak{J}_s(c^*, w_1) + 2 = \mathfrak{J}_{i,s}((b^*, c^*), w_1)$$

So (b^*, c^*) does not dominate (b, c) .

- If $(b^*, c^*) \in DS$, then, by (2),

$$\mathfrak{J}_{i,s}((b, c), w_1) = s(1, t) + s(0, 1 - t + \varepsilon) + 1 - t$$

$$< \mathfrak{s}(1, 1-t) + \mathfrak{s}(0, 1-t) + 2 - t < \mathfrak{J}_{i,\mathfrak{s}}((b^*, c^*), w_1)$$

So (b^*, c^*) does not dominate (b, c) .

- If $(b^*, c^*) \in SB$, then, by (3),

$$\begin{aligned} \mathfrak{J}_{i,\mathfrak{s}}((b, c), w_1) &= \mathfrak{s}(1, t) + \mathfrak{s}(0, 1-t+\varepsilon) + 1 - t \\ &< \mathfrak{s}(1, t) + \mathfrak{s}(0, t) + 2 - t < \mathfrak{J}_{i,\mathfrak{s}}((b^*, c^*), w_1) \end{aligned}$$

So (b^*, c^*) does not dominate (b, c) .

- If $(b^*, c^*) \in DD$, then, by (4) and (A),

$$\begin{aligned} \mathfrak{J}_{i,\mathfrak{s}}((b, c), w_1) &= \mathfrak{s}(1, t) + \mathfrak{s}(0, 1-t+\varepsilon) + 1 - t \\ &< \mathfrak{s}(1, 1-t) + 1 \leq \mathfrak{J}_{i,\mathfrak{s}}((b^*, c^*), w_1) \end{aligned}$$

So (b^*, c^*) does not dominate (b, c) .

- If $(b^*, c^*) \in BB$, then, by (5) and (B),

$$\begin{aligned} \mathfrak{J}_{i,\mathfrak{s}}((b, c), w_1) &= \mathfrak{s}(1, t) + \mathfrak{s}(0, 1-t+\varepsilon) + 1 - t \\ &< \mathfrak{s}(0, t) + 1 \leq \mathfrak{J}_{i,\mathfrak{s}}((b^*, c^*), w_1) \end{aligned}$$

So (b^*, c^*) does not dominate (b, c) .

- If $(b^*, c^*) \in SS$, then, by (6),

$$\begin{aligned} \mathfrak{J}_{i,\mathfrak{s}}((b, c), w_1) &= \mathfrak{s}(1, t) + \mathfrak{s}(0, 1-t+\varepsilon) + 1 - t \\ &< \mathfrak{s}(1, t) + \mathfrak{s}(0, 1-t) + 2 - 2t < \mathfrak{J}_{i,\mathfrak{s}}((b^*, c^*), w_1) \end{aligned}$$

So (b^*, c^*) does not dominate (b, c) .

- If $(b^*, c^*) \in BS$ accuracy dominates (b, c) relative to $\mathfrak{J}_{i,\mathfrak{s}}$, then, by (7), c^* accuracy dominates c relative to \mathfrak{s} . But, since $\mathfrak{J}_{\mathfrak{s}}$ is connected, we know that all credence functions that accuracy dominate c are in SS . So (b^*, c^*) does not accuracy dominate (b, c) .
- If $(b^*, c^*) \in SD$ accuracy dominates (b, c) relative to $\mathfrak{J}_{i,\mathfrak{s}}$, then, by (8), c^* accuracy dominates c relative to \mathfrak{s} . But, since $\mathfrak{J}_{\mathfrak{s}}$ is connected, we know that all credence functions that accuracy dominate c are

in SS . So (b^*, c^*) does not accuracy dominate (b, c) .

- If $(b^*, c^*) \in BD$, then, by (9),

$$\begin{aligned} \mathfrak{J}_{i,\mathfrak{s}}((b, c), w_2) &= \mathfrak{J}_{\mathfrak{s}}(c, w_2) + 1 + 1 - t \\ &< \mathfrak{J}_{\mathfrak{s}}(c^*, w_2) + 2 = \mathfrak{J}_{i,\mathfrak{s}}((b^*, c^*), w_2) \end{aligned}$$

So (b^*, c^*) does not dominate (b, c) .

Thus, there is $\varepsilon > 0$ such that the above inequalities hold. For that ε , we have that $((B, S), (t, 1-t+\varepsilon))$ is undominated.

This result allows us to prove our two theorems:

- *Proof of Theorem 3* It is straightforward to calculate that, if we let $t = 0.7$ and $1-t+\varepsilon = 0.6$, the above inequalities hold, and thus $(b_{\text{Cleo}}, c_{\text{Cleo}}) = ((B, S), (0.7, 0.6))$ is undominated. \square

- *Proof of Theorem 4* Note that, for all $t > 0.5$,

$$\begin{aligned} -q(1, t) + q(0, 1-t) - q(1, 1-t) &= 2(1-t)^2 - t^2 < t \\ -q(1, t) + q(0, 1-t) - q(0, t) &= 2(1-t)^2 - t^2 < t \end{aligned}$$

Then let $b = (B, S)$ and $c = (t, 1-t+\varepsilon)$. \square

- *Proof of Theorem 5* If \mathfrak{s} is a normalised strictly proper scoring rule that generates a connected inaccuracy measure $\mathfrak{J}_{\mathfrak{s}}$. Then we simply pick r such that

$$\begin{aligned} -\mathfrak{s}(1, r) + \mathfrak{s}(0, 1-r) - \mathfrak{s}(1, 1-r) &< t \\ -\mathfrak{s}(1, r) + \mathfrak{s}(0, 1-r) - \mathfrak{s}(0, r) &< t \end{aligned}$$

As noted above, this is always possible. Moreover, for $t > r$, these inequalities continue to hold. Then let $b = (B, S)$ and $c = (t, 1-t+\varepsilon)$. \square

This completes our proofs. \square

References

- Arló-Costa, H. and Pedersen, A. P. (2012). Belief and Probability: A General Theory of Probability Cores. *International Journal of Approximate*

- mate Reasoning*, 53(3):293–315.
- Easwaran, K. (ms). Dr. Truthlove, or, How I Learned to Stop Worrying and Love Bayesian Probability.
- Easwaran, K. and Fitelson, B. (2012). An 'Evidentialist' Worry about Joyce's Argument for Probabilism. *Dialectica*, 66(3):425–433.
- Fitelson, B. (ms). *Coherence*. Oxford University Press, Oxford, UK.
- Foley, R. (1992). The Epistemology of Belief and the Epistemology of Degrees of Belief. *American Philosophical Quarterly*, 29(2):111–124.
- Foot, P. (1967). The Problem of Abortion and the Doctrine of Double Effect. *Oxford Review*, 5:5–15.
- Hempel, C. (1962). Deductive-Nomological vs. Statistical Explanation. In Feigl, H. and Maxwell, G., editors, *Minnesota Studies in the Philosophy of Science*, volume III, pages 98–169. University of Minnesota Press, Minneapolis.
- Joyce, J. M. (1998). A Nonpragmatic Vindication of Probabilism. *Philosophy of Science*, 65(4):575–603.
- Joyce, J. M. (2009). Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief. In Huber, F. and Schmidt-Petri, C., editors, *Degrees of Belief*. Springer.
- Kyburg, H. E. (1961). *Probability and the Logic of Rational Belief*. Wesleyan University Press, Connecticut.
- Leitgeb, H. (2014). The Stability Theory of Belief. *Philosophical Review*, 123(2):131–171.
- Levinstein, B. A. (2012). Leitgeb and Pettigrew on Accuracy and Updating. *Philosophy of Science*, 79(3):413–424.
- Pettigrew, R. (2013). Accuracy and Evidence. *Dialectica*, 67(4):579–96.
- Rosenkrantz, R. D. (1981). *Foundations and Applications of Inductive Probability*. Ridgeview Press, Atascadero, CA.
- Thomson, J. J. (1976). Killing, Letting Die, and the Trolley Problem. *The Monist*, 59(2):204–217.