# Accuracy and Infinity: A Dilemma for Subjective Bayesians*

Mikayla Kelley[†] and Sven Neth[‡]

[†]Stanford University (mikelley@stanford.edu)
[‡]University of California, Berkeley (nethsven@berkeley.edu)

Penultimate Draft. Forthcoming in *Synthese*

**Abstract**

We argue that subjective Bayesians face a dilemma: they must offend against the spirit of their permissivism about rational credence or reject the principle that one should avoid accuracy dominance.

## 1 Introduction

According to subjective Bayesianism, one's credences prior to receiving evidence are rationally permissible so long as they satisfy the axioms of probability. We show that subjective Bayesians face a dilemma: they must offend against the spirit of their permissivism about rational credence or they must reject the principle that one should avoid accuracy dominance. This is surprising, because a popular way to motivate subjective Bayesianism is to appeal to an equivalence between satisfying the axioms of probability and avoiding accuracy dominance. However, this equivalence only holds when an agent has finitely many credences, which gives rise to our dilemma. We spell out the dilemma in detail and then consider how the subjective Bayesian might respond.

*We would like to thank Wesley Holliday, Mathias Böhm, Kiran Luecke, Snow Zhang and two anonymous referees for helpful comments on earlier drafts which have made this paper significantly better.

## 2 The Dilemma

According to Bayesian epistemologists, agents have degrees of belief or *credences* in *propositions*. For example, an agent might have credence .5 that it is going to rain tomorrow or credence .25 that a fair coin will land heads twice in a row. We model propositions as subsets of a set $W$ of *possible worlds*, and we model one's credences with a *credence function* $c : \mathcal{F} \to [0,1]$ for $\mathcal{F} \subseteq \mathcal{P}(W)$ an *opinion set* of propositions. Here, $p \in \mathcal{F}$ represents the fact that the agent assigns some credence to $p$, and $c(p)$ represents the agent's credence in $p$.

According to *subjective* Bayesians, one's initial credence function prior to receiving evidence relevant to the opinion set—what we will call one's *prior* credence function—is rationally permissible insofar as it is *coherent*, that is, satisfies the axioms of probability.[1] So ascribing to subjective Bayesianism involves, by definition, accepting the following principle:

> **Subjective Credence**: A prior credence function is rationally permissible if and only if it is coherent.

*Objective* Bayesians, on the other hand, deny **Subjective Credence**; they think that there are constraints on one's prior credence function beyond coherence, such as the Principle of Indifference (Joyce 2004, pp. 423-31). These constraints might single out a unique permissible prior credence function or a class of permissible prior credence functions (Carnap 1952). We pose our dilemma for subjective Bayesians since it most sharply targets their view, but the dilemma applies to plausible versions of objective Bayesianism as well (see Footnote 10 for further discussion).

So according to the subjective Bayesian, besides the minimal constraint of coherence, it is 'up to you' which prior credences to adopt. From a subjective

[1] More precisely, $c : \mathcal{F} \to [0,1]$ is coherent just in case it can be extended to a finitely additive probability function—there is a function $\overline{c} : \mathcal{A}(\mathcal{F}) \to [0,1]$, where $\mathcal{A}(\mathcal{F})$ is the algebra generated by $\mathcal{F}$, with the following properties:

- $\overline{c}(W) = 1$;
- $\overline{c}(A \cup B) = \overline{c}(A) + \overline{c}(B)$ for $A, B \in \mathcal{A}(\mathcal{F})$ with $A \cap B = \varnothing$;
- $\overline{c}(A) = c(A)$ for all $A \in \mathcal{F}$.

One might ask why we do not assume that a credence function is defined on an algebra. First, there is a robust tradition within the accuracy literature where credences are defined over arbitrary (finite) opinion sets (e.g., Pettigrew 2016a; Predd et al. 2009); second, we will consider a principle which explicitly allows agents to have an opinion set other than an algebra, so we would not want to rule out such a principle by the definition of a credence function; finally, as we will discuss in detail later, the dilemma we pose goes through even if credence functions are defined on algebras.

Bayesian point of view, it then looks natural to think *which propositions you assign prior credences to* is 'up to you' in the same way. That is, it looks like the following principle is very much in the spirit of subjective Bayesianism:

> **Subjective Opinion Set**: One is rationally permitted to assign prior credences to any set of propositions.

Up until recently, it seemed that **Subjective Credence** and **Subjective Opinion Set** are not only consistent with but motivated by concern for *accuracy*. In particular, it is well known that on any finite opinion set, coherence is equivalent to avoiding *accuracy dominance*, where credence function $c$ accuracy dominates credence function $d$ if $d$ is less accurate than $c$ in all possible worlds (de Finetti 1974; Joyce 1998; Schervish et al. 2009; Predd et al. 2009; Pettigrew 2016a).[2] This equivalence between coherence and avoiding accuracy dominance motivates both the requirement of coherence and the thought that *nothing more* than coherence is needed for rational credences. So subjective Bayesians have reason to accept the following principle:[3]

> **Avoid Accuracy Dominance**: If a credence function is accuracy dominated, then it is not rationally permissible.

Our dilemma arises for the subjective Bayesian in light of work by Kelley (forthcoming), which shows that the equivalence between coherence and avoiding accuracy dominance breaks down when an agent has infinitely many credences.[4] On certain infinite opinion sets, coherent credence functions can be accuracy dominated. Here is an example from Kelley forthcoming:

> (Example 1) Let $\mathcal{F} = \{\{n \leq N : n \in \mathbb{N}\} : N \in \mathbb{N}\}$ be an opinion set over $\mathbb{N}$ (including zero), and let $c(\{n \leq N\}) = 0$ for each $N \in \mathbb{N}$.

[2]More precisely, we measure inaccuracy with a function $\mathscr{I} : \mathcal{C} \times W \to [0, \infty]$, where $\mathcal{C}$ is the set of credence functions on an opinion set $\mathcal{F} \subseteq \mathcal{P}(W)$. We say that $c$ *strongly* dominates $d$ if $\mathscr{I}(d, w) > \mathscr{I}(c, w)$ for all $w \in W$. We say that $c$ *weakly* dominates $d$ if $\mathscr{I}(d, w) \geq \mathscr{I}(c, w)$ for all $w \in W$ and $\mathscr{I}(d, w) > \mathscr{I}(c, w)$ for some $w \in W$. Throughout, we will not specify whether we are speaking of weak or strong dominance unless the distinction matters.

[3]As Joyce (1998, pp. 580-4) points out, de Finetti, one of the major proponents of subjective Bayesianism, laid the groundwork for accuracy dominance arguments for coherence and so would presumably be committed to this principle.

[4]While we draw on Kelley's formal results at various points in our argument, we part from her work in exploring what the failure of the equivalence between avoiding accuracy dominance and coherence entails for subjective Bayesians. She, on the other hand, focuses on offering as general of an accuracy dominance argument as possible for Probabilism, the principle that one ought to have (at least) coherent credences.

Intuitively, this credence function might model one's prior credences about a countably infinite fair lottery. You think that a positive integer will be drawn at random; since the draw is random and from an infinite set, you think that the probability that any particular positive integer will be drawn is zero; by finite additivity, you also think that the probability that the drawn number is less than or equal to any particular positive integer is zero.

Now, note that $c$ is coherent, and so according to **Subjective Credence** and **Subjective Opinion Set**, $c$ is rationally permissible. However, $c$ has infinite inaccuracy in all possible worlds and is therefore accuracy dominated. To define accuracy, we start by assuming a popular accuracy measure called the *Brier* score (Brier 1950) but relax this assumption in Section 3. Where the *omniscient credence function* $v_w$ at a possible world $w \in W$ is the coherent credence function which assigns 1 to $p$ if $p$ contains $w$ and 0 otherwise, the Brier score of credence function $c$ at $w$ is the 'distance', in some sense, between $c$ and $v_w$. More precisely, the Brier score of $c$ at $w$ is $\sum_{p\in\mathcal{F}}(v_w(p) - c(p))^2$. Since, in the example, the Brier score of $c$ at any world is an infinite sum of 1s and an omniscient credence function has finite inaccuracy at every world, $c$ is accuracy dominated by an omniscient credence function. So according to **Avoid Accuracy Dominance**, $c$ is *not* rationally permissible. This shows that **Subjective Credence**, **Subjective Opinion Set**, and **Avoid Accuracy Dominance** are inconsistent. Therefore, the subjective Bayesian must either reject **Subjective Opinion Set** or reject **Avoid Accuracy Dominance**. This is our dilemma.

In response to the example above, one might say that when considering credences over countably infinite opinion sets, *countable additivity*—and not just finite additivity—should be imposed.[5] This would rule out the credence function above, because it is not countably additive. Countable additivity is controversial, with many subjective Bayesians rejecting it as a requirement of rationality—in particular de Finetti and Savage (Lyon 2016; Liu 2020). But if countable additivity were to secure the avoidance of accuracy dominance in the infinite setting, this would perhaps show that subjective Bayesians should accept countable additivity after all.

However, Kelley presents another example which shows that countable additivity is not enough to avoid accuracy dominance:

[5] A credence function $c : \mathcal{F} \to [0, 1]$ is countably additive if $c(A) = \sum_{i=1}^{\infty} c(A_i)$ for $\{A_i\}_{i=1}^{\infty} \subseteq \mathcal{F}$ with $A_i \cap A_j = \varnothing$ for $i \neq j$.

(Example 2) Let $\mathcal{F} = \{\{n \geq N : n \in \mathbb{N}\} : N \in \mathbb{N}\}$ be an opinion set over $\mathbb{N}$ (including zero), and let $c(\{n \geq N\}) = \frac{1}{\sqrt{N+1}}$ for each $N \in \mathbb{N}$.

Intuitively, this credence function might model one's prior credences about some physical process. The collision of two quantum particles will emit some random amount of energy $n$, and you have credences about how much energy will be emitted. We may even suppose that these credences match our best physical theory.

Note that $c$ is coherent and extendable to a countably additive probability function. However, $c$ has infinite inaccuracy in all possible worlds and is therefore accuracy dominated.[6] So strengthening the notion of coherence underlying **Subjective Credence** does not mitigate our dilemma.

Another initial thought is to weaken **Avoid Accuracy Dominance**: as Pettigrew (2016a) has argued, perhaps what is rationally impermissible is to have a credence function which is dominated *by an undominated credence function*. However, this response also will not help since in both examples presented, $c$ is dominated by an undominated credence function. In both cases, $c$ is dominated by an omniscient credence function, and omniscient credence functions are never accuracy dominated.

In the next section, we consider more promising responses to our dilemma.

# 3 Responses

In response to our dilemma, the subjective Bayesian might reject **Subjective Opinion Set**. Kelley (forthcoming) introduces a number of constraints on an infinite opinion set that restore the equivalence between coherence and avoiding accuracy dominance. For example, if one has credences on a countably infinite partition, then coherence is equivalent to avoiding accuracy dominance (Theorem 4.26).

However, we think that requiring rational agents to have credences over particular kinds of opinion sets is hard to square with the motivations behind subjective Bayesianism. At the heart of subjective Bayesianism is a commitment to significant freedom in adopting priors. We can think of priors as encoding

[6]The Brier score of $c$ at world $N \in \mathbb{N}$ is bounded below by $\sum_{n>N} \frac{1}{n+1}$ which is infinite (since the harmonic series is divergent), so any omniscient credence function (strongly) dominates $c$, as an omniscient credence function is finitely inaccurate at each world.

the evidence-independent epistemic tendencies that determine how an agent responds to evidence—what one might call an agent's *epistemic standards*.[7] These include one's sensitivity to epistemic risk, whether one tends to prioritize avoiding falsities or seeking truths, how extreme one is in one's opinions, and one's degree of skepticism. We can capture different responses to the same body of total evidence by reference to differences in these epistemic standards.

Now, the subjective Bayesian thinks that vastly different responses to the same body of total evidence can be rationally permissible: so long as these epistemic standards can be encoded into a probability distribution, they are rationally permitted. But notice that an agent's epistemic standards will *also* affect the contents of their opinion set. For example, a more risk averse epistemic agent might be more hesitant to have any credence at all in a particular proposition. One's opinion set also reflects what one takes to be a genuine possibility for the purpose of reasoning and decision-making, which depends on one's epistemic standards. For example, when I assign credences to a coin landing heads and tails but don't include the possibility that the coin lands on its edge, this might be because, in light of my less open-minded epistemic standards, I don't take this possibility seriously. Thus, given that subjective Bayesians are permissive regarding epistemic standards and epistemic standards also influence the propositions one bothers to form credences about, subjective Bayesians have reason to be permissive regarding rational opinion sets. Of course, there is no logical incoherence in accepting subjective Bayesianism and rejecting **Subjective Opinion Set**. But, given what we have said, it seems unlikely that this position could be motivated. Permissivism about rational credences naturally leads to permissivism about rational opinion sets.[8]

One might insist that some *minimal* constraint on the structure of the opinion set would be in the spirit of subjective Bayesianism, for such a minimal constraint would be the analogue to the subjective Bayesian's insistence that prior credences be at least coherent. However, given that this constraint on the opinion set would need to be minimal to preserve the spirit of subjective

[7] For discussion of this interpretation of prior credences, see Schoenfield 2014, Meacham 2016, and Titelbaum 2022, Sec. 4.3.2.

[8] This discussion raises the more general questions of whether there are any rationality constraints on an agent's opinion set or the underlying space of possibilities, what such rationality constraints would look like, and how they relate to more standard rationality constraints on credence functions such as, e.g., the Principle of Indifference. In particular, can objective Bayesians motivate rationality constraints on the opinion set and the underlying space of possibilities? Steele and Stefánsson (2021) discuss the related issue of how to model unawareness and conceptual change in a Bayesian framework. For more related discussion, see Carr 2015; Talbot 2019; Hewson 2021.

Bayesianism, we conjecture that the dilemma would remain. For example, the most natural constraint to impose would be that the opinion set form an algebra,[9] but there are coherent credence functions on algebras that are accuracy dominated. Indeed, when extended to algebras, the credence functions from Examples 1 and 2 are (weakly) dominated by omniscient credence functions. In fact, on the algebra generated by $\mathcal{F}$ in both examples, the *only* coherent credence functions not (weakly) dominated are the omniscient credence functions (see the Appendix for proof). By this result, it also will not help if one assumes that a credence function is both defined on an algebra and countably additive. We will return to this result soon, but for now note that it shows that requiring one's opinion set to be an algebra is not going to resolve our dilemma.

Another way to reject **Subjective Opinion Set** would be to require that an agent have credences in only finitely many propositions. However, this not only clashes with the permissive spirit of subjective Bayesianism but is also in tension with two important components of subjective Bayesianism: *long-run convergence theorems*, which show that under certain conditions, different priors 'wash out' when updated on the same evidence, and *representation theorems*, which derive subjective probabilities from preferences. Subjective Bayesians frequently appeal to long-run convergence theorems to argue that, despite having different priors, rational agents will agree in the long run (Earman 1992, Ch. 6). These convergence theorems usually require credences over infinitely many propositions. Representation theorems, on the other hand, derive subjective probabilities from preferences and are often seen as providing a foundation for the notion of subjective probability. Many representation theorems, in particular the representation theorem of Savage (1972), require credences over infinitely many propositions. So in light of the way these two important components of subjective Bayesianism rely on infinite opinion sets, insistence on finitely many credences does not look open to the subjective Bayesian.

A third way to reject **Subjective Opinion Set** is to require that the opinion set be a partition. As noted, Kelley establishes that if the opinion set is a countably infinite partition, the equivalence between coherence and avoiding accuracy dominance is restored in the countably infinite setting. However, we do not see any reason to think that it would be *impermissible* for one's opinion set to be an algebra, which is what one is committed to when choosing this path

[9] Although Lyon (2016), drawing on Fine 1973, discusses reasons to be skeptical that one's opinion set must always be an algebra.

to reject **Subjective Opinion Set**.[10]

Turn now to the second horn of the dilemma: rejecting **Avoid Accuracy Dominance**. There are different ways of doing so. First, one might think that it is sometimes okay to adopt accuracy dominated credences according to an acceptable notion of accuracy dominance. We think that this is very implausible. If a credence function $d$ is less accurate than a credence function $c$ in all possible worlds relative to an acceptable measure of inaccuracy, this is a conclusive reason to think that $d$ is not permissible, especially if $c$ is not itself dominated.[11]

Second, one might restrict **Avoid Accuracy Dominance** to the finite setting, perhaps because there is no acceptable notion of accuracy dominance in the infinite setting. This would be a way of escaping our dilemma without permitting accuracy dominated credences. However, we think that it would be bad news for subjective Bayesians if there was simply *no* accuracy dominance argument in the infinite setting. Indeed, one of the central motivations for subjective Bayesianism is that on any (finite) opinion set, one does as well as one can do in terms of avoiding accuracy dominance just in case one is coherent. If this fact is restricted to finite opinion sets, then this central motivation is not fully general: the subjective Bayesian cannot point to accuracy as a reason why prior credences on infinite opinion sets should be coherent and why mere coherence suffices. The subjective Bayesian might, of course, appeal to some other justification, for example pragmatic dutch book arguments, which apply to infinite opinion sets (for a detailed discussion, see Nielsen 2020). Still, the fact remains that considerations of accuracy—a core epistemic virtue—would not provide a fully general vindication of subjective Bayesianism. Moreover, as discussed before, credence functions on infinite opinion sets are crucial to a number of important ideas in subjective Bayesianism and so cannot be ignored.

[10] We think that our dilemma also poses problems for objective Bayesians because it seems that the credence function in Example 2 is rationally permissible—after all, we can suppose that these credences match our best physical theory. However, there are more responses available to the objective Bayesian than to the subjective Bayesian. For example, objective Bayesians might respond to our dilemma by rejecting **Subjective Opinion Set** and appealing to constraints on one's opinion set, for example, demanding that opinion sets be finite or form a partition. Objective Bayesians might also argue that the coherent credence functions in Examples 1 and 2 are not rationally permissible by defending constraints on rational prior credences. We think that such constraints will be hard to defend, because in many cases, the *only* credence functions which avoid accuracy dominance are omniscient credence functions (see Appendix); and it is implausible to think that rationality demands that we adopt an omniscient credence function.

[11] Pettigrew (2016a), Chapter 2 provides a thorough discussion of dominance reasoning in the accuracy-first setting.

Thus, we think that, rather than permitting accuracy dominated credences or restricting **Avoid Accuracy Dominance** to the finite setting, subjective Bayesians have very good reason to look for a way of spelling out **Avoid Accuracy Dominance** that is i) consistent with **Subjective Opinion Set** and ii) rules out incoherent credence functions without ruling out coherent credence functions in the infinite setting. So the subjective Bayesian should suggest a reformulation of the notion of accuracy dominance assumed thus far. As we will discuss now, there are formidable challenges in doing so. To see this, let us consider three implicit assumptions underlying **Avoid Accuracy Dominance**, as we understand it, that the subjective Bayesian might reject.

First, we have been assuming that the inaccuracy of a credence function on a countably infinite opinion set is measured with a natural extension of the Brier score. Many have argued that the Brier score is a legitimate method of scoring inaccuracy (Horwich 1982, Maher 2002, Joyce 2009, Pettigrew 2016a); some have even argued that it is the *unique* legitimate method of scoring inaccuracy (Leitgeb and Pettigrew 2010). Still, one might deny that we should measure inaccuracy using the Brier score, in general or just in the infinite setting.

Pruss (forthcoming) proves a number of relevant impossibility results when the opinion set is uncountable. Here we show that there are problems for this strategy even in the countably infinite setting. Indeed, there are two kinds of measures that we are aware of relative to which avoiding accuracy dominance is equivalent to coherence in the countably infinite setting: the inaccuracy measure suggested by Walsh (ms) and the 'admissible' inaccuracy measures considered by Nielsen (forthcoming). Both give up on *additivity*, where an inaccuracy measure is additive if, roughly, the inaccuracies of the individual credences contribute equally to the inaccuracy of the full credence function.[12] Many accuracy dominance arguments that the subjective Bayesian might appeal to in the finite case assume additivity (e.g., Predd et al. 2009, Leitgeb and Pettigrew 2010, Pettigrew 2016b).[13] Walsh's inaccuracy measure weights the contributions of each individual credence to the overall inaccuracy of the credence function such that the weights can be ordered and tend to zero. Nielsen assumes that admissible

[12]See the definition of *quasi-additivity* in the Appendix. Quasi-additivity covers a wide variety of inaccuracy measures besides the Brier score, including all the measures considered by Predd et al. (2009). For example, the *log score*, an inaccuracy measure defended by McCutcheon 2019 among others, is quasi-additive. The log score of a credence function $c$ at a world $w$ is given by $\sum_{p\in\mathcal{F}} -\ln|1 - v_w(p) - c(p)|$.

[13]Though see Pettigrew 2022 (see also Nielsen 2022). As our discussion shows, this work of Pettigrew's may prove very important to the subjective Bayesian in responding to our dilemma.

inaccuracy measures have a number of features, but additivity—or something even approximating additivity—is not among them.

In fact, we can say something more general here: *any* inaccuracy measure according to which coherence is equivalent to avoiding accuracy dominance in the countably infinite setting must give up on additivity (see the Appendix for a formal result). Thus, Nielsen's admissible inaccuracy measures *cannot* be additive or even approximately additive. More precisely, we show that on a wide class of countably infinite opinion sets, including many algebras, if accuracy dominance is defined according to an (approximately) additive inaccuracy measure, then a coherent credence function will avoid accuracy dominance if and only if it is omniscient.[14]

Moreover, for a large range of opinion sets, additive inaccuracy measures will fail to be what Nielsen calls *quasi-strictly proper*. Many accuracy dominance arguments that the subjective Bayesian might appeal to in the finite case assume the stronger condition of *strict propriety*. The idea behind propriety conditions is that from the perspective of a coherent credence function, no other credence function should be more accurate on expectation. In the finite setting, additivity and (quasi-)strict propriety are perfectly consistent, but this is not true in the infinite setting. As our and Nielsen's results together show, keeping the weaker quasi-strict propriety (as well as making a number of other assumptions) but dropping additivity leads to a fully general accuracy dominance argument for subjective Bayesianism, while keeping additivity but giving up quasi-strict propriety entails that there is no fully general accuracy dominance argument for subjective Bayesianism. We conclude that it is not easy to formulate a version of **Avoid Accuracy Dominance** that will motivate subjective Bayesianism in the infinite setting by merely changing the measure of inaccuracy—rather, doing so requires returning to the difficult and unresolved question of which conditions must be met by a plausible measure of inaccuracy.

A second assumption underlying **Avoid Accuracy Dominance** is that an agent's entire credence function is used for the purpose of assigning inaccuracy scores. However, one might instead assign inaccuracy scores to a credence function at a world by scoring only a representative part or transformation of the

[14] This result assumes dominance is understood as *weak* dominance. As we prove in the Appendix, no coherent credence function is strongly dominated on the class of countably infinite opinion sets we consider, but no *incoherent* credence function is strongly dominated either. Thus, moving to a version of **Avoid Accuracy Dominance** in terms of strong dominance while keeping additivity is not going to help the subjective Bayesian motivate their view in the infinite setting, since any such principle will not differentiate between coherence and incoherence on a wide class of countably infinite opinion sets, including many algebras.

credence function at that world. Let us call such a way of scoring inaccuracy *representative scoring*.[15] To see the potential problem with using representative scoring to offer an accuracy dominance argument for subjective Bayesianism, let's work with a particular instance of the method.

Assume that credence function $c$ is defined on an opinion set over a countably infinite set of worlds which contains all the singletons, and let $c|_W$ be $c$ restricted to the singleton sets. For example, extending $c$ from Example 2 in the natural way to be defined on the singletons, we have:

$$c|_W(\{N\}) = c(\{n \geq N\}) - c(\{n \geq N+1\}) = \frac{1}{\sqrt{N+1}} - \frac{1}{\sqrt{N+2}}$$

for each $N \in \mathbb{N}$. One might argue that an agent's epistemic state can be entirely captured, for the purposes of determining inaccuracy, by its credences that each world is the actual world. One could then define the inaccuracy of $c$ at $w$ to be, for example, the Brier score of $c|_W$ at $w$. As noted above, Kelley proves that on countably infinite partitions, a credence function is dominated (using the Brier score) if and only if it is incoherent. So it follows that if $c$ is coherent, then $c|_W$ is coherent and so undominated.

However, if $c$ is *incoherent*, it does not follow that $c|_W$ is incoherent, for the agent's incoherence might arise with respect to more complex propositions, for example the disjunctive proposition represented by $\{w_1, w_2\}$. Therefore, on this instance of representative scoring, the incoherence of $c$ does *not* imply that $c$ is accuracy dominated. Thus, the subjective Bayesian loses the ability to offer an accuracy dominance argument for having (at least) coherent credences.

More generally, the problem with representative scoring is that there is no guarantee that if a credence function is incoherent, then so is its representative for the purpose of scoring inaccuracy. Put another way, it is possible that an incoherent and a coherent credence function are assigned the same inaccuracy scores using representative scoring, which makes it impossible to offer an accuracy argument for **Subjective Credence**. Of course, there may be some other way of taking this route that avoids this problem, but the (arguably) most natural version—scoring a credence function by scoring its restriction to the singletons—will not work. Thus, we leave it as a challenge to the subjective Bayesian to offer a representative scoring method that enables an accuracy dominance argument for **Subjective Credence**.

A third implicit assumption of **Avoid Accuracy Dominance** as we have

[15]Thanks to an anonymous referee for suggesting this objection.

stated it is that *global* accuracy dominance should be avoided, not just what we might call *local* accuracy dominance. More precisely, **Avoid Accuracy Dominance** assumes that dominance relations between credence functions are established by comparing the inaccuracies assigned to each credence function at each world. This is how accuracy dominance is usually understood in the finite setting.[16]

However, one might instead suggest that such global dominance relations do not matter; what matters is whether the *finite parts* of the credence function are accuracy dominated. Say that a credence function is *locally accuracy dominated* if the restriction of the credence function to some finite subset of the full opinion set is accuracy dominated.[17] A result of Schervish et al. (2009) shows that a credence function on an opinion set of any size is coherent if and only if it is not locally accuracy dominated. This result, along with Examples 1 and 2, show that coherent credence functions can be *globally* accuracy dominated without being *locally* accuracy dominated. The subjective Bayesian could then argue that it is local accuracy dominance which ought to be avoided, not global accuracy dominance.

This response faces its own challenges. First, it is far from clear why avoiding local dominance is the only thing that matters. To think that the rationality of one's entire epistemic state supervenes on the rationality of its proper parts seems to be an instance of the fallacy of composition. Presumably an epistemic state can be irrational while having proper parts which, taken alone, are rational. Moreover, even if the rationality of one's epistemic state *did* supervene on its proper parts, we see no reason why rationality would supervene on only its *finite* proper parts.[18] But once we broaden the scope of local accuracy dominance and take into account some infinite parts as well, the dilemma posed here will likely kick in: there will be coherent credence functions that are locally accuracy dominated in this sense. So if the subjective Bayesian were to go this route in responding to our dilemma, they face the following challenge: they must come up with a plausible formulation of a local accuracy dominance principle according to which coherence is equivalent to avoiding accuracy dominance even in the infinite setting.

[16] See Kelley forthcoming, Sec. 6.1 for discussion of this point.

[17] More precisely, let $c : \mathcal{F} \to [0, 1]$ be *locally accuracy dominated* if $c|_{\mathcal{X}}$ is accuracy dominated relative to the Brier score for some finite subset $\mathcal{X} \subseteq \mathcal{F}$.

[18] Thanks to Snow Zhang for this point.

## 4 Conclusion

We have shown that subjective Bayesians face a difficult choice: accept constraints on which propositions one can have credences in—constraints which are at odds with the permissive spirit of their view—or reject the principle that one should avoid accuracy dominance. We think that the best way out of the dilemma is to come up with a plausible accuracy dominance principle that can be used to motivate subjective Bayesianism in both the finite and infinite settings. As we have shown, even this response faces some formidable challenges.

## 5 Appendix

**Definition 5.1.** Let an opinion set $\mathcal{F}$ on $W$ be *rich* if $\mathcal{F}$ includes all finite subsets of $W$.

**Definition 5.2.** Assume $\mathcal{F}$ is countable. Let an inaccuracy measure $\mathscr{I}$ be *quasi-additive* if it has the form

$$\mathscr{I}(c, w) = \sum_{p \in \mathcal{F}} \lambda_p \mathfrak{d}(v_w(p), c(p)),$$

where $\{\lambda_p\}_{p \in \mathcal{F}} \subseteq (0, \infty)$ and $\mathfrak{d} : \{0, 1\} \times [0, 1] \to [0, \infty]$ is a function such that $\mathfrak{d}(x, y) = 0$ if and only if $x = y$.

We take it that a necessary condition on $\mathscr{I}$ being at least approximately additive is that $\sum_{p \in \mathcal{G}} \lambda_p = \infty$ for any infinite subset $\mathcal{G} \subseteq \mathcal{F}$.

**Definition 5.3.** For $c$ a coherent credence function on $\mathcal{F} \subseteq \mathcal{P}(W)$, let $c^*$ be a *probabilistic extension* of $c$ to $\mathcal{P}(W)$ if $c^*$ is a finitely additive probability function on $\mathcal{P}(W)$ with $c^*(W) = 1$ and $c^*(p) = c(p)$ for all $p \in \mathcal{F}$.[19] Let an inaccuracy measure $\mathscr{I}$ be *quasi-strictly proper* if $\mathbb{E}_{c^*}\mathscr{I}(c, \cdot) \leq \mathbb{E}_{c^*}\mathscr{I}(d, \cdot)$ for all coherent credence functions $c$, probabilistic extensions $c^*$ of $c$, and credence functions $d$ on $\mathcal{F}$, with strict inequality if $d$ is incoherent.[20]

**Theorem 5.4.** Let $\mathcal{F}$ be a rich, countably infinite opinion set.[21] Let $\mathscr{I}$ be a

[19]By the definition of coherence, $c$ is extendable to a finitely additive probability function $\bar{c}$ on $\mathcal{A}(\mathcal{F})$, the algebra generated by $\mathcal{F}$. By Corollary 3.3.4 from Rao and Rao 1983, $\bar{c}$ is extendable to a finitely additive probability function $c^*$ on $\mathcal{P}(W)$. Note that the proof of Corollary 3.3.4 uses a result that is equivalent to the Axiom of Choice.

[20]See Nielsen forthcoming for a relevant discussion of taking expectations with a merely finitely additive probability function.

[21]See Pruss (forthcoming) for a related impossibility result in the uncountable case (Proposition 3).

quasi-additive inaccuracy measure with weights such that $\sum_{p \in \mathcal{G}} \lambda_p = \infty$ for any infinite subset $\mathcal{G} \subseteq \mathcal{F}$. Then:

(1) no credence function on $\mathcal{F}$ is strongly dominated relative to $\mathscr{I}$;

(2) a credence function $c$ on $\mathcal{F}$ is not weakly dominated relative to $\mathscr{I}$ if and only if $c$ is an omniscient credence function;

(3) $\mathscr{I}$ is not quasi-strictly proper.

*Proof.* We start by proving that under the assumptions of the theorem, (*) any credence function is finitely inaccurate at no more than one world. We then prove (1) and (2) from (*).

Assume a credence function $c$ is finitely inaccurate at some world $w \in W$. Then fix $v \neq w$ and consider a sequence $\{r_i^v\}_{i=1}^{\infty}$ of elements in $\mathcal{F}$ with $w \in r_i^v$ and $v \notin r_i^v$ for each $i$ (such a sequence exists by the richness of $\mathcal{F}$). The inaccuracy of $c$ at $w$ is bounded below by $\sum_{i=1}^{\infty} \lambda_{r_i^v} \mathfrak{d}(1, c(r_i^v))$. Thus, $c(r_i^v) \to 1$ as $i \to \infty$. Indeed, if $c(r_i^v)$ did not tend to 1, then one could chose a subsequence of $\{c(r_i^v)\}_{i=1}^{\infty}$ that is bounded away from 1; since $\mathscr{I}$ is quasi-additive and we have assumed that infinite sets of weights have infinite sum, $c$ would then be infinitely inaccurate at $w$. Now the inaccuracy of $c$ at $v$ is bounded below by $\sum_{i=1}^{\infty} \lambda_{r_i^v} \mathfrak{d}(0, c(r_i^v))$. Since $c(r_i^v) \to 1$, it follows by quasi-additivity and the assumption that infinite sets of weights have infinite sum that $\sum_{i=1}^{\infty} \lambda_{r_i^v} \mathfrak{d}(0, c(r_i^v))$ is infinite. So $c$ is infinitely inaccurate at $v$. Thus, every credence function is finitely inaccurate at no more than one world.

(1) immediately follows from (*): two credence functions which are both infinitely inaccurate at some world cannot strongly dominate one another.

As for (2), to see that omniscient credence functions are not weakly dominated, note that if an omniscient credence function $v_w$ were weakly dominated by some credence function $c$, then $\mathscr{I}(c, w) = 0$ (since $\mathscr{I}(v_w, w) = 0$). Since $\mathscr{I}$ is quasi-additive, if $\mathscr{I}(c, w) = 0$, then $c = v_w$; but clearly a credence function cannot weakly dominate itself. So $v_w$ is not weakly dominated relative to $\mathscr{I}$. For the other direction, if a credence function $c \neq v_w$ is infinitely inaccurate at every world, then it is weakly dominated by an omniscient credence function. If $c$ is not infinitely inaccurate at every world, then by (*) it is finitely inaccurate at a single world, say, $w$. It follows that $c$ is weakly dominated by $v_w$: $c$ is infinitely inaccurate at all worlds $v \neq w$; and at $w$, since $v_w \neq c$, it follows by quasi-additivity that $\mathscr{I}(c, w) \neq 0$ while $\mathscr{I}(v_w, w) = 0$.

As for (3), we will show that there are coherent credence functions on $\mathcal{F}$ which are infinitely inaccurate at every world. By the definition of a finitely additive expectation of an extended-real function (see, e.g., the appendix of Nielsen forthcoming), $\mathbb{E}_{c^*}\mathscr{I}(c,\cdot) = \infty$ for any probabilistic extension $c^*$ of $c$ when $c$ has infinite inaccuracy at every world; and so $\mathbb{E}_{c^*}\mathscr{I}(c,\cdot)$ is not strictly less than $\mathbb{E}_{c^*}\mathscr{I}(d,\cdot)$ for an incoherent credence function $d$. Therefore, $\mathscr{I}$ is not quasi-strictly proper.

As an example of a coherent credence function which is infinitely inaccurate at every world, consider $c = \frac{1}{2}v_{w_1} + \frac{1}{2}v_{w_2}$ for $w_1 \neq w_2$. It is easy to check that $c$ is coherent. To see that it is infinitely inaccurate at every world, consider the following sequences of elements in $\mathcal{F}$ (which can be found by the richness of $\mathcal{F}$):

(a) $\{r_i^{w_1}\}_{i=1}^{\infty}$ such that $w_1 \in r_i^{w_1}$ and $w_2 \notin r_i^{w_1}$ for each $i$;

(b) $\{r_i^{w_2}\}_{i=1}^{\infty}$ such that $w_2 \in r_i^{w_2}$ and $w_1 \notin r_i^{w_2}$ for each $i$;

(c) for each $v \neq w_1, w_2$, let $\{r_i^{v}\}_{i=1}^{\infty}$ be such that $v \in r_i^{v}$ and $w_1, w_2 \notin r_i^{v}$ for each $i$.

The inaccuracy of $c$ at $w_1$ is bounded below by $\sum_{i=1}^{\infty} \lambda_{r_i^{w_1}} \mathfrak{d}(v_{w_1}(r_i^{w_1}), c(r_i^{w_1})) = \sum_{i=1}^{\infty} \lambda_{r_i^{w_1}} \mathfrak{d}(1, \frac{1}{2})$, which is infinite by quasi-additivity and the assumption that infinite collections of weights have infinite sum. Analogous reasoning establishes that $c$ is infinitely inaccurate at $w_2$. For $v \neq w_1, w_2$, note that $\mathscr{I}(c, v)$ is bounded below by $\sum_{i=1}^{\infty} \lambda_{r_i^{v}} \mathfrak{d}(v_v(r_i^{v}), c(r_i^{v})) = \sum_{i=1}^{\infty} \lambda_{r_i^{v}} \mathfrak{d}(1, 0)$, which is infinite by quasi-additivity and the assumption that infinite collections of weights have infinite sum. Thus, $c = \frac{1}{2}v_{w_1} + \frac{1}{2}v_{w_2}$ is infinitely inaccurate at every world, establishing that $\mathscr{I}$ is not quasi-additive. $\square$

## References


Glenn W Brier. Verification of forecasts expressed in terms of probability. *Monthly Weather Review*, 78(1):1–3, 1950. doi:10.1175/1520-0493(1950)078<0001:VOFEIT>2.0.CO;2.

Rudolf Carnap. *The Continuum of Inductive Methods.* Chicago, IL, USA: University of Chicago Press, 1952.

Jennifer Rose Carr. Epistemic expansions. *Res Philosophica*, 92(2):217–236, 2015. doi:10.11612/resphil.2015.92.2.4.

Bruno de Finetti. *Theory of Probability*. John Wiley, New York, 1974.

John Earman. *Bayes or Bust? A Critical Examination of Bayesian Confirmation Theory*. MIT Press, 1992.

Terrence L. Fine. *Theories of Probability: An Examination of Foundations*. Academic Press, 1973.

Matthew Hewson. Accurate believers are deductively cogent. *Noûs*, 2021. doi:10.1111/nous.12397.

Paul Horwich. *Probability and Evidence*. Cambridge University Press, 1982.

James M. Joyce. A nonpragmatic vindication of probabilism. *Philosophy of Science*, 65(4):575–603, 1998. doi:10.1086/392661.

James M. Joyce. The development of subjective bayesianism. In Dov M. Gabbay, John Woods, and Akihiro Kanamori, editors, *Handbook of the History of Logic*, pages 10–415. Elsevier, 2004. doi:10.1016/b978-0-444-52936-7.50012-4.

James M. Joyce. Accuracy and coherence: Prospects for an alethic epistemology of partial belief. In Franz Huber and Christoph Schmidt-Petri, editors, *Degrees of Belief*, pages 263–297. Springer, 2009. doi:10.1007/978-1-4020-9198-8_11.

Mikayla Kelley. On accuracy and coherence with infinite opinion sets. *Philosophy of Science*, forthcoming. doi:10.1017/psa.2021.37.

Hannes Leitgeb and Richard Pettigrew. An objective justification of Bayesianism I: Measuring inaccuracy. *Philosophy of Science*, 77(2):201–235, 2010. doi:10.1086/651317.

Yang Liu. Countable additivity, idealization, and conceptual realism. *Economics and Philosophy*, 36(1):127–147, 2020. doi:10.1017/s0266267118000536.

Aidan Lyon. Kolmogorov's axiomatisation and its discontents. *The Oxford Handbook of Probability and Philosophy*, pages 155–166, 2016. doi:10.1093/oxfordhb/9780199607617.013.9.

Patrick Maher. Joyce's argument for probabilism. *Philosophy of Science*, 69(1): 73–81, 2002. doi:10.1086/338941.

Randall G. McCutcheon. In favor of logarithmic scoring. *Philosophy of Science*, 86(2):286–303, 2019. doi:10.1086/702028.

Christopher J. G. Meacham. Ur-priors, conditionalization, and ur-prior conditionalization. *Ergo: An Open Access Journal of Philosophy*, 3, 2016. doi:10.3998/ergo.12405314.0003.017.

Michael Nielsen. The strength of de finetti's coherence theorem. *Synthese*, 198 (12):11713–11724, 2020. doi:10.1007/s11229-020-02825-7.

Michael Nielsen. On the best accuracy arguments for probabilism. *Philosophy of Science*, 89(3):621–630, 2022. doi:10.1017/psa.2021.43.

Michael Nielsen. Accuracy and probabilism in infinite domains. *Mind*, forthcoming.

Richard Pettigrew. *Accuracy and the Laws of Credence*. Oxford University Press UK, 2016a. doi:10.1093/acprof:oso/9780198732716.001.0001.

Richard Pettigrew. Accuracy, risk, and the principle of indifference. *Philosophy and Phenomenological Research*, 92(1):35–59, 2016b. doi:10.1111/phpr.12097.

Richard Pettigrew. Accuracy-first epistemology without additivity. *Philosophy of Science*, 89(1):128–151, 2022. doi:10.1017/psa.2021.17.

Joel Predd, Robert Seiringer, Elliott Lieb, Daniel Osherson, H. Vincent Poor, and Sanjeev Kulkarni. Probabilistic coherence and proper scoring rules. *IEEE Transactions on Information Theory*, 55(10):4786–4792, 2009. doi:10.1109/tit.2009.2027573.

Alexander Pruss. Accuracy, probabilism and bayesian update in infinite domains. *Synthese*, forthcoming. doi:10.1007/s11229-022-03938-x.

K.P.S Bhaskara Rao and M. Bhaskara Rao. *Theory of Charges: A Study of Finitely Additive Measures*. Academic Press Inc., 1983.

Leonard J. Savage. *The Foundations of Statistics*. Wiley Publications in Statistics, second revised edition edition, 1972.

Mark Schervish, Teddy Seidenfeld, and Joseph B. Kadane. Proper scoring rules, dominated forecasts, and coherence. *Decision Analysis*, 6(4):202–221, 2009. doi:10.1287/deca.1090.0153.

Miriam Schoenfield. Permission to believe: Why permissivism is true and what it tells us about irrelevant influences on belief. *Noûs*, 48(2):193–218, 2014. doi:10.1111/nous.12006.

Katie Steele and H. Orri Stefánsson. *Beyond Uncertainty: Reasoning with Unknown Possibilities*. Cambridge University Press, 2021. doi:10.1017/9781108582230.

Brian Talbot. Repugnant accuracy. *Noûs*, 53(3):540–563, 2019. doi:10.1111/nous.12237.

Michael Titelbaum. *Fundamentals of Bayesian Epistemology 1: Introducing Credences*. Oxford University Press, 2022. doi:10.1093/oso/9780198707608.001.0001.

Sean Walsh. Probabilism in infinite dimensions. ms. Unpublished manuscript.