A Contextual Accuracy Dominance Argument for Probabilism*

Mikayla Kelley

University of Chicago

Draft of July 2023

Abstract

A central motivation for Probabilism—the principle of rationality that requires one to have credences that satisfy the axioms of probability—is the accuracy dominance argument: one should not have accuracy dominated credences, and one avoids accuracy dominance just in case one satisfies Probabilism. Until recently, the accuracy dominance argument for Probabilism has been restricted to finite credal sets. One reason for this is that it is not clear how to measure the accuracy of infinitely many credences in a motivated way. In particular, as recent work has shown, the conditions often imposed in the finite setting are mutually inconsistent in certain infinite settings. In this paper, I offer a fully general accuracy dominance argument for Probabilism that is consistent with that not all sets of credences can be measured for accuracy. The normative core of the argument is the principle that one should not have credences that would be accuracy dominated in some epistemic context one might find oneself in if there are alternative credences which do not have this defect. An important upshot of the proposed argument is a general "contextualizing" strategy for extending accuracy arguments restricted to the finite setting.

1. Introduction

A core tenant of Bayesian epistemology is that one's degrees of belief or *credences* ought to satisfy the axioms of probability. Call this principle *Probabilism.* There is a long tradition going back to de Finetti (1974) of defending Probabilism by showing that it promotes accuracy.[1] Accuracy arguments for Probabilism have three crucial ingredients: necessary conditions on a "legitimate" measure of inaccuracy, a normative accuracy-rationality bridge principle in terms of a legitimate inaccuracy measure, and a mathematical theorem that establishes that one's credences are ruled out by the normative principle if one's credences violate Probabilism (Pettigrew 2022).

*For helpful discussion, thanks to Kevin Dorst, Dan Friedman, Wesley Holliday, Sven Neth, Teddy Seidenfeld, Jer Steeger, Francesca Zaffora-Blando, Snow Zhang and the audience at the Pittsburgh Formal Epistemology Workshop and the Formal Epistemology Workshop 2023.

[1]See, e.g., Joyce 1998, Joyce 2004, Schervish et al. 2009, Predd et al. 2009, Leitgeb and Pettigrew 2010a, Leitgeb and Pettigrew 2010b, Pettigrew 2016a, Nielsen 2023, and Kelley 2023.

One such argument that has been especially influential is the following version of the accuracy dominance argument, which I will soon present in more detail. A legitimate inaccuracy measure is continuous, additive, and strictly proper; if a credence function is accuracy dominated by another credence function which is itself not accuracy dominated, then it is rationally impermissible; a credence function is accuracy dominated by an undominated credence function if and only if it fails to satisfy Probabilism; so, Probabilism.

While influential, this argument is restricted in a crucial way: it only works to justify Probabilism when the credal set is finite. This is a serious limitation. First, rationality does not obviously preclude credence functions with infinite domains. Indeed, there are important components of Bayesian epistemology which assume rational agents have such credence functions.[2] Moreover, as Nielsen (2023) points out, probability functions on infinite domains are ubiquitous in the sciences, and presumably it is sometimes rationally permissible to match one's credences to those recommended by the sciences. Second, it seems plausible that we could actually have infinitely many credences once we recognize that we could do so without having infinitely many mental states. In particular, we could have credences that make use of a variable ranging over some infinite set; for example, I might have credence $(\frac{1}{2})^n$ that a fair coin would land hands n times in a row if flipped n times in a row.[3] Finally, Bayesianism is notoriously known as a theory of ideal rationality. So, even if finite agents like us cannot in fact have infinitely many credences, this would not entail that Probabilism as a principle of rationality is restricted to finite opinion sets, for Probabilism—like much of Bayesian epistemology—is arguably best viewed as a rationality constraint on cognitively ideal agents. So there is good reason to seek an extension of the accuracy dominance argument for Probabilism to the infinite setting.

However, there are formidable obstacles that stand in the way of doing so.[4] As we will see, additive, continuous, and strictly proper inaccuracy measures do not exist in general for infinite sets of credences. One response to this fact is to consider weaker conditions on a legitimate inaccuracy measure than those often assumed in the finite case (Nielsen 2023). But as I will argue, this response risks relying on illegitimate inaccuracy measures in one's accuracy argument for Probabilism; and an accuracy argument is only as convincing as the measures of inaccuracy used are plausible.

Another response takes seriously the existing arguments for what kinds of features a legitimate inaccuracy measure should have and concedes the possibility that there is simply no legitimate way to measure inaccuracy for certain sets of credences. One then defends a normative accuracy-rationality bridge principle which does not require that there be legitimate ways to measure inaccuracy in full generality. If the math works out, one would then have a fully general accuracy argument for Probabilism. This is precisely the route I take here.

To give a bit more high-level introduction, the accuracy-rationality bridge principle that I defend relies on the idea that for almost every question we face, not every proposition we have a credence

[2] Convergence theorems, which show that under certain conditions, differences in priors 'wash out' when updated on the same evidence, often make use of infinite opinion sets (Gaifman and Snir 1982; Schervish and Seidenfeld 1990; Earman 1992, Ch. 6). Also, a number of representation theorems, which derive subjective probabilities from preferences, require credences over infinitely many propositions (Savage 1972).

[3] See Pruss 2022 for further discussion of this point.

[4] For thorough discussion of the problems faced in extending the accuracy dominance argument, see Pruss 2022 and Kelley and Neth 2023.

in will be relevant. Thus, it is natural to consider how accurate credal states are not just holistically but also relative to particular questions or *epistemic contexts* in which accuracy is well-defined. I suggest the following norm on such contextualized credences: one should not have credences that are accuracy dominated in some epistemic context if there are alternative credences which do not have this defect. I defend this norm and argue that it offers an accuracy-based justification for Probabilism no matter the size of one's credal state. Moreover, it does so without having to resort to making use of potentially illegitimate measures of inaccuracy.

The plan for the paper is as follows. In Section 2, I review the standard accuracy dominance argument for Probabilism. In Section 3, I review some impossibility results. In Section 4, I criticially discuss Nielsen's (2023) recent response to these impossibility results and motivate the value of a new kind of response. In Section 5, I present such a response, which I call the *Contextual Accuracy Dominance Argument for Probabilism.* Section 6 briefly covers some questions for future research. I conclude in Section 7.

2. The Finite Global Accuracy Dominance Argument for Probabilism

In this section, I present the standard accuracy dominance argument for Probabilism used in the finite setting. We begin by reviewing the formal framework. We represent *propositions* using subsets of a set W of possible worlds, and we represent an agent's credences with a credence function $c : \mathcal{F} \to [0, 1]$ for $\mathcal{F} \subseteq 2^W$ an *opinion set* of propositions. On this model, $p \in \mathcal{F}$ represents the fact that the agent assigns some credence to p, and $c(p)$ represents the agent's credence in p. A credence function is *coherent* just in case it satisfies the axioms of probability and *incoherent* if not.[5] Let $\mathcal{C}_\mathcal{F}$ denote the set of credence functions on $\mathcal{F}$ and $\mathcal{P}_\mathcal{F}$ denote the set of coherent credence functions on $\mathcal{F}$. An important set of coherent credence functions are the *omniscient* credence functions: for $w \in W$, the omniscient credence function at w is the function v_w which takes value 1 if $w \in p$ and 0 if $w \notin p$. *Probabilism* is then the rationality principle that incoherent credence functions are rationally impermissible.

The standard accuracy dominance argument for Probabilism starts with a formal representation of the inaccuracy of credences on some $\mathcal{F} \subseteq 2^W$.[6] This representation is a function that takes in a credence function c on $\mathcal{F}$ and a possible world $w \in W$ and outputs $\mathscr{I}(c, w) \in [0, \infty]$, the inaccuracy of the credence function at that world. With an inaccuracy measure in hand, we can define the notion of *accuracy dominance*: given two credence functions c and c',

- c *strongly dominates* c' relative to $\mathscr{I}$ if $\mathscr{I}(c', w) > \mathscr{I}(c, w)$ for all worlds w;

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

[6] It has become standard to work with inaccuracy measures rather than accuracy measures; this is merely a convention.

- c *weakly dominates* c' relative to $\mathscr{I}$ if $\mathscr{I}(c,' w) \geq \mathscr{I}(c, w)$ for all worlds w with strict inequality for some world.

Less formally, one credence function accuracy dominates another if it is less inaccurate no matter how the world turns out to be.

One of the central components of an accuracy argument is a claim about what the necessary conditions are on a *legitimate* inaccuracy measure. Here are a number of conditions that will be relevant to our discussion. For some of these conditions, we will refer to expected inaccuracy, which requires the notion of a *probabilistic extension* of a coherent credence function: c^* defined on 2^W is a probabilistic extension of c on $\mathcal{F} \subseteq \mathcal{P}(W)$ if $c^* = c$ on $\mathcal{F}$ and c^* is a finitely additive probability function.[7]

- Continuity: For each $w \in W$, $\mathscr{I}(c, w)$ is a continuous function of c on $\mathcal{C}_{\mathcal{F}}$;[8]

- Weak Continuity: For each $w \in W$, $\mathscr{I}(c, w)$ is a continuous function of c on $\mathcal{P}_{\mathcal{F}}$;

- Weak Continuity in Expectation: For every $p \in \mathcal{P}_{\mathcal{F}}$ and probabilistic extension p^* of p, $\mathbb{E}_{p^*}\mathscr{I}(c, \cdot)$ is a continuous function on $\mathcal{P}_{\mathcal{F}}$;

- Additivity: For $\mathcal{F}$ countable, $\mathscr{I}(c, w) = \sum_{p \in \mathcal{F}} \lambda_p \mathfrak{d}(v_w(p), c(p))$ for weights $\{\lambda_p\}_{p \in \mathcal{F}} \subseteq (0, \infty)$ and some $\mathfrak{d} : \{0, 1\} \times [0, 1] \to [0, \infty]$;

- Strong Additivity: For $\mathcal{F}$ countable, $\mathscr{I}(c, w) = \sum_{p \in \mathcal{F}} \lambda_p \mathfrak{d}(v_w(p), c(p))$ for weights $\{\lambda_p\}_{p \in \mathcal{F}} \subseteq (0, \infty)$ such that $\sum_{p \in \mathcal{G}} \lambda_p = \infty$ for every infinite collection $\mathcal{G} \subseteq \mathcal{F}$ and some $\mathfrak{d} : \{0, 1\} \times [0, 1] \to [0, \infty]$ such that $\mathfrak{d}(x, y) = 0$ if and only if $x = y$;

- Strict Propriety: $\mathbb{E}_{p^*}\mathscr{I}(p, \cdot) < \mathbb{E}_{p^*}\mathscr{I}(c, \cdot)$ for all coherent credence functions p, some probabilistic extension p^* of p, and all credence functions $c \neq p$ on $\mathcal{F}$;

- Quasi-Strict Propriety: $\mathbb{E}_{p^*}\mathscr{I}(p, \cdot) \leq \mathbb{E}_{p^*}\mathscr{I}(c, \cdot)$ for all coherent credence functions p on $\mathcal{F}$, some probabilistic extension p^* of p, and all credence functions c on $\mathcal{F}$, with strict inequality if c is incoherent;[9]

- Local boundedness: For every $p \in \mathcal{P}_{\mathcal{F}}$, $\mathscr{I}(p, w)$ is bounded on W.

There are a number of accuracy arguments for various rationality principles that make different assumptions about the necessary conditions on a legitimate inaccuracy measure (see Pettigrew 2022 for a helpful overview of these different assumptions). But as Pettigrew discusses, the best—or at least a popular—version of the accuracy dominance argument for Probabilism assumes that a legitimate

[7]One could define expected inaccuracy using extensions of coherent credence functions to smaller algebras. However, since every coherent credence function can be extended to a finitely additive probability function on the full power set (by, e.g., Theorem 3.2.10 of Rao and Rao 1983), we will avoid dealing with technical issues by assuming extensions are defined on the full power set.

[8]When we move to the infinite setting, what we mean by continuity here and in the next two definitions becomes less clear. We will discuss this in detail later on.

[9]Later on, we will look at a stronger version of quasi-strict propriety on which the relevant inequalities hold for *all* probabilistic extensions of a coherent credence function.

inaccuracy measure is additive, continuous, and strictly proper. So this is the first premise of the standard accuracy dominance argument for Probabilism as I will understand it here.

The second premise is an accuracy-rationality bridge principle in terms of accuracy dominance, which I will call *Undominated Dominance* following Pettigrew (2016a):

> **Undominated Dominance** (UD): If a credence function is strongly accuracy dominated relative to a legitimate inaccuracy measure by a credence function which is not even weakly dominated relative to a legitimate inaccuracy measure, then it is rationally impermissible.[10]

The third premise is a mathematical theorem establishing that relative to additive, continuous, and strictly proper inaccuracy measures, an incoherent credence function on a finite opinion set is strongly dominated by a coherent credence function and no coherent credence function is even weakly dominated.

Here then is a first pass at the standard accuracy dominance argument for Probabilism.

P1) Any legitimate inaccuracy measure is additive, continuous, and strictly proper.

P2) Undominated Dominance

P3) (Predd et al. 2009) Relative to any additive, continuous, and strictly proper inaccuracy measure, an incoherent credence function is strongly dominated by a coherent credence function, and coherent credence functions are not even weakly dominated.

C) So, Probabilism

As has already been hinted at, this presentation of the standard accuracy dominance argument contains false advertising: the mathematical result in (P3) only holds for finite opinion sets, and so Probabilism has only been defended for credence functions on finite opinion sets despite being a principle that applies to all credence functions. But there is another problem with this presentation of the argument. When applying Undominated Dominance, there is the presumption that a legitimate inaccuracy measure exists for every opinion set on which Probabilism is to be defended via the argument. For if a legitimate inaccuracy measure did not exist on all such opinion sets, then the accuracy dominance relations referenced in Undominated Dominance would sometimes not be well-defined. This hidden assumption in the standard argument will prove to be crucial in generalizing the argument to the infinite setting.

With these points taken into account, here is the final statement of the standard accuracy dominance argument for a restricted version of Probabilism. For reasons that will become clear, I call it the *Finite Global Accuracy Dominance Argument for Probabilism.*

Finite Global Accuracy Dominance Argument for Probabilism (FGAP)

[10] There is a stronger principle that one might use instead: if a credence function is strongly dominated, then it is rationally impermissible. However, as Pettigrew convincingly argues, this principle is too strong in light of cases where every credence function is dominated by some other credence function. Presumably, in such cases, not every credence function is impermissible.

P1) Any legitimate inaccuracy measure is additive, continuous, and strictly proper.

P2) Undominated Dominance

P3) (Predd et al. 2009) Let $\mathcal{F}$ be a finite opinion set. Relative to any additive, continuous, and strictly proper inaccuracy measure on $\mathcal{F}$, an incoherent credence function on $\mathcal{F}$ is strongly dominated by a coherent credence function on $\mathcal{F}$, and coherent credence functions on $\mathcal{F}$ are not even weakly dominated.

P4) There exists a legitimate inaccuracy measure for any finite opinion set.

C) So, if one has finitely many credences, one ought to satisfy Probabilism.

It is this argument whose generalization to the infinite setting we seek.

3. Impossibility Results

As discussed in the previous section, three commonly assumed features of a legitimate inaccuracy measure are additivity, continuity, and strict propriety. Such conditions are consistent on finite opinion sets, but this is not so on infinite opinion sets in general.

First, there are no strictly proper inaccuracy measures on uncountably infinite power sets.[11]

Theorem 1 (Pruss 2022)**.** If $\mathcal{F}$ is an uncountable power set, then there is no strictly proper inaccuracy measure on $\mathcal{F}$.[12]

The proof is enlightening: on such uncountable domains, there are more coherent credence functions than assigments of inaccuracy values to worlds; and so two coherent credence functions will have the same inaccuracy at every world, which is inconsistent with strict propriety. Thus, we cannot generally impose strict propriety in the uncountably infinite setting because it requires that we can assign inaccuracy values in a more fine-grained way than is possible with real-valued inaccuracy measures.

This raises the question of whether we might formally represent inaccuracy with something other than a real-valued function. For example, we might consider functions taking values in something other than the reals (see Pruss 2022, Sec. 4.2 for discussion); or we might work with qualitative accuracy orderings of credence functions (though Zhang (ms) proves a very general impossibility result in this setting).[13] I will set these possibilities aside for now. I hope to show by the end of the essay that we can stick with real-valued functions and still successfully defend Probabilism in full generality.

Turn now to additivity. A first problem is that it is far from obvious how to generalize additivity to the uncountable setting. Kelley (2023, Sec. 5) considers measuring inaccuracy using integration against a kind of uniform measure. However, such inaccuracy measures cannot offer a vindication of Probabilism but only a weaker principle that says, roughly, one should have *almost* coherent credences. Second, there is an impossibility result even for countably infinite opinion sets if we assume strong

[11] As far as I know, it is an open question whether Theorem 1 extends to all uncountable opinion sets.

[12] Beyond this result, Pruss (2022) offers an in-depth investigation into various impossibility results that arise when trying to measure inaccuracy in the infinite setting.

[13] See also Park and Jung 2023.

additivity. In the presence of continuity and strict propriety in the finite setting, strong additivity and additivity are equivalent.[14] Thus, we could see (P1) of FGAP as requiring that a legitimate inaccuracy measure be strongly additive.

But relative to strongly additive inaccuracy measures on a wide range of countably infinite opinion sets, no credence function is strongly accuracy dominated and only omniscient credence functions are not weakly dominated. Moreover, on such countably infinite opinion sets, there are no strongly additive *and* quasi-strictly proper inaccuracy measures.

Definition 2. Let an opinion set $\mathcal{F}$ on W be *rich* if $\mathcal{F}$ contains all finite subsets W.

Theorem 3 (Kelley and Neth 2023)**.** Let $\mathcal{F}$ be a rich, countably infinite opinion set. Let $\mathscr{I}$ be a strongly additive inaccuracy measure. 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.

The upshot: we cannot keep both (P1) and (P4) in extending FGAP to infinite opinion sets. For there to exist a legitimate inaccuracy measure for any opinion set, the conditions on a legitimate inaccuracy measure must be weakened from those often imposed in the finite setting.

4. Weakening the Necessary Conditions on a Legitimate Inaccuracy Measure

One response to these impossibility theorems is thus to impose weaker conditions on a legitimate inaccuracy measure than those often imposed in the finite case. This is what Nielsen (2023) does and proves a theorem that can be used to defend a very general accuracy dominance argument for Probabilism.

Theorem 4 (Nielsen 2023)**.** Assume c is defined on 2^W for some set W of worlds. Let $\mathscr{I}$ be weakly continuous in expectation,[15] quasi-strictly proper, and locally bounded. Then if c is incoherent, it is strongly dominated relative to $\mathscr{I}$ by a coherent credence function, and if c is coherent, then it is not even weakly dominated relative to $\mathscr{I}$.

Here is the associated general accuracy dominance argument for Probabilism:

<u>Nielsen's Global Accuracy Dominance Argument for Probabilism (NGAP)</u>

N1) Any legitimate inaccuracy measure is weakly continuous in expectation, quasi-strictly proper, and locally bounded.

[14] This follows from Theorem 4.3.5 of Pettigrew 2016a, which offers a characterization of additive and continuous strictly proper inaccuracy measures (on finite opinion sets) in terms of what are called *Bregman divergences*.

[15] Nielsen assumes the product topology on $\mathcal{C}_{2^W}$. As we will see, the second half of this result can be proved with a weaker notion of continuity. We will also prove the second half of this result for credence functions on arbitrary opinion sets, not just power sets of an underlying set of worlds.

N2) Undominated Dominance

N3) Theorem 4

N4) There exists a legitimate inaccuracy measure for any opinion set.[16]

C) So, Probabilism.

NGAP, as stated, is also false advertising to a certain extent, since Theorem 4 assumes that credence functions are defined on a power set of a set of worlds. Thus, Probabilism has only been defended for credence functions on power sets. Nielsen mentions without proof that the argument could be extended to credence functions on arbitrary algebras (see his Footnote 13). Still, it is not clear that rationality requires having credences on an algebra.[17] It may be that Theorem 4 could be extended to credence functions on arbitrary opinion sets, but this is yet to be determined (though see the proof of Theorem 7 for a partial extension). I will set this issue aside for now to raise a different, I think, more pressing worry for NGAP.

If it is true that *only* weak continuity in expectation, quasi-strict propriety, and local boundedness are needed to be a legitimate inaccuracy measure, then NGAP offers a promising extension of FGAP to the infinite setting. But if, in fact, a legitimate inaccuracy measure should have these three features *and more*—e.g., strict propriety or strong additivity—then we should start to worry that sometimes no such measures exist. In that case, (N4) would be false, and NGAP would be unsound.

I am not going to try to systematically argue that the conditions that NGAP imposes on a legitimate inaccuracy measure are too weak—that, in fact, something like strong additivity or the full strength of strict propriety is required for a legitimate measure of inaccuracy. A closer examination of the conditions assumed in Nielsen's argument would be worthwhile, but that is not what I will offer here. Instead, I am going to try to argue for the weaker claim that it would be of value to develop a new kind of accuracy dominance argument which does not weaken the conditions on a legitimate inaccuracy measure from those assumed in (P1) and does not assume that every opinion set admits of a legitimate inaccuracy measure. This is all that is needed to motivate my proposed alternative generalization of FGAP. To this end, I have four things to say.

First, there is an extensive literature on the necessary conditions on a legitimate inaccuracy measure, and many have argued for necessary conditions that are in some way stronger than those imposed in NGAP; in particular, they impose constraints that are sometimes impossible to meet in the infinite setting. A natural—and arguably the most consistent—response to this impossibility is to simply admit that legitimate inaccuracy measures do not always exist in the infinite setting. Taking on board NGAP instead puts one in a somewhat difficult dialectical position: one must reject these arguments for the conditions on a legitimate inaccuracy measure offered in the finite setting or explain why those arguments only apply if the opinion set is finite. It must be admitted that it would be at least

[16] Technically, he only needs the weaker assumption that there exists a legitimate inaccuracy measure on any power set.

[17] See, e.g., Fine 1973 and Lyon 2016. See also discussion of the rationality principle *Subjective Opinion Set* in Kelley and Neth 2023.

convenient if one could avoid either of these options and still have a fully general accuracy dominance argument for Probabilism.

Second, I will offer a source of skepticism about weakening strict propriety to quasi-strict propriety in a motivated way.[18] For a full defense of NGAP, one needs an argument for quasi-strict propriety. But this argument cannot also be an argument for strict propriety, since all such an argument would show is that (N4) is false, as not all opinion sets would admit of legitimate inaccuracy measures. Moreover, the only existing arguments for propriety of any kind are for *weak propriety*—a condition weaker than quasi-strict propriety—and for strict propriety.[19] So there seems to be good reason to be at least unsure at this point that there exists an argument for the precise strength of the propriety condition needed for NGAP.

Third, strong additivity is an appealing property of inaccuracy measures, one that formalizes four commonsense ideas about inaccuracy. Commitment to strong additivity is, in effect, a commitment to the following four ideas:

a) the inaccuracy of an overall credal state is a weighted aggregation of the inaccuracies of individual credences;

b) how the inaccuracy of an individual credence is scored does not depend on the content of the proposition the credence is about;

c) there is a non-zero, minimum importance placed on the inaccuracy of each credence to the overall inaccuracy of the credal state;[20]

d) each individual credence is assigned the perfect accuracy score if and only if it is the ideal credence, that is, assigns 1 to truths and 0 to falsities.

(a)-(c) capture something like *minimal egalitarianism for credal states*: in scoring the accuracy of an overall credal state, individual credences should be treated roughly the same and given some minimal importance. (d) is necessary to formalize a basic thought underwriting the accuracy framework: the inaccuracy of a credence is, in some sense, its distance from the ideally accurate credence, where the ideally accurate credence is 1 if the proposition is true and 0 if the proposition is false. Strong additivity thus formalizes a number of natural thoughts about inaccuracy and so would be costly to give up.

Finally, even if one is fully convinced that NGAP succeeds as an accuracy argument for Probabilism, there is still value in having a second accuracy argument on the table. All else equal, two good arguments for a philosophical position are better than one.

In light of these considerations, I suggest that there is value in looking for another kind of accuracy dominance argument for Probabilism, one that does not potentially rely on illegitimate measures of inaccuracy. This is what I offer in the next section.

[18] Pruss (2022) makes a similar point.

[19] See Campbell-Moore and Levinstein 2021.

[20] It is easy to see that the condition on strong additivity that every infinite set of weights have infinite sum is equivalent to the weights having positive infimum.

Before moving on though, let me consider an objection. One might worry that the possible inability to measure the inaccuracy of every credence function is itself too much to admit. Such restrictedness might be seen as in itself a serious limitation of accuracy-first epistemology.[21] However, I think it is too early to conclude this. The success of accuracy-first epistemology is not determined by to what extent plausible formal representations of inaccuracy exist per se but rather to what extent epistemic principles of rationality can be vindicated in their full generality using plausible formal representations of inaccuracy. And we might be able to accomplish the latter to a significant extent while having plausible formal representations of inaccuracy in only a restricted range of cases. Indeed, the contextual accuracy dominance argument to come in the next section is evidence that we can vindicate Probabilism in its full generality while having plausible formal representations of inaccuracy in only a restricted range of cases.

5. A Contextual Accuracy Dominance Argument for Probabilism

In this section, I will present the Contextual Accuracy Dominance Argument for Probabilism. We begin by noticing that to answer any given question, not all of one's credences will be relevant. When trying to determine the cost of a flight to Mexico, one's credences regarding the number of squirrels in your backyard is irrelevant; when inquiring into whether coffee is healthier than tea, one's credence in the Goldbach conjecture is irrelevant; when inquiring into the weight of the moon, one's credence in moral realism is irrelevant; and so forth. Thus, while we have the notion of an agent's entire credal state, represented by a credence function on an opinion set, we should also have the notion of the part of the agent's credal state which is relevant for answering a particular question.[22] Thus, let an *epistemic context* (for the agent) be a subset of the agent's opinion set unified in their being the propositions in the agent's opinion set which are relevant to answering some question; and for c a credence function on an opinion set $\mathcal{F}$ and $\mathcal{S} \subseteq \mathcal{F}$, let $c|_{\mathcal{S}}$ represent the part of the agent's credal state relevant for answering the question (or questions) associated with the epistemic context $\mathcal{S}$.[23]

In adopting this formal representation of epistemic contexts, I am making two assumptions. First, I am assuming that every set of propositions constitutes an epistemic context and thus are the propositions relevant to answering some question an agent could face. Second, I am assuming that for every question, there is a principled way to delimit precisely those propositions relevant to answering that question. With regards to the first assumption, I find it plausible that for every set of propositions, we could cook up some question for which the propositions in that set are exactly the relevant ones. However, I will not spend much time defending this assumption because it could be weakened and the contextual accuracy dominance argument on offer would remain sound. Indeed, one could require

[21] Thanks to Francesca Zaffora Blando for this objection.

[22] In related work, Horowitz (2019) suggests that credences are accurate insofar as they license true educated guesses. Here, we highlight the importance of accuracy with respect to particular questions but stick with the orthodox idea that accuracy (with respect to a question) is distance to the relevant truths.

[23] The formal results to follow could be adopted with a different philosophical interpretation of subsets of an agent's opinion set. For example, we might think of subsets as representing propositions relevant not just to answering some question but to more generally completing some task (such as planning, deciding, acting, etc.) Or one might think of (some of) these subsets as representing the propositions one would have non-zero credence in upon learning some piece of evidence. Thanks to Weng Kin San and Zach Barnett for helping me here.

epistemic contexts to be not just any set of propositions but a set of propositions that meets some minimal condition C, so long as C itself meets the following criterion: C must be such that for every finite set of propositions X, there is some finite set of propositions X' which contains X and satisfies C. For this reason, I will not spend much time responding to the objection that not every set of propositions constitutes an epistemic context for an agent. With regards to the second assumption, I recognize that this is an idealizing assumption and leave relaxing it to future work.

Now, given that we (almost) never find ourselves answering questions for which our full credal state is relevant, it is puzzling why most existing accuracy dominance arguments focus only on scoring full credal states. It would also seem to matter—even from a purely epistemic perspective—how accurate a credal state is within an epistemic context. In particular, insofar as having a credal state which is accuracy dominated is rationally impermissible if there are credal states available which are not so dominated, it seems plausible that having a credal state which is accuracy dominated within some epistemic context is rationally impermissible insofar as there are credal states which are not so dominated. Moreover, the rational failure of being accuracy dominated in a context is an epistemic, non-pragmatic failure just like being accuracy dominated full stop; one's credences are failing in their functional role of accurately representing the facts (albeit a more restricted set of facts than one has views on), not necessarily in their functional role within planning and decision-making. In this way, we can turn to focus on accuracy dominance within a context without giving up on the non-pragmatic character of the accuracy dominance argument for Probabilism which has always been a large part of its appeal.[24]

But which contexts should be quantified over in a contextual dominance principle of rationality? If all of them, we will face the impossibility results discussed in Section 3. If not all of them, we need a principled way of constraining the contexts in which one ought not be dominated. One initial thought is to include only finite contexts. The relevant contextual dominance principle would then say: do not have a credence function which, when restricted to a finite subset, is accuracy dominated by a credence function which is itself undominated in all finite contexts. Then via a result of Schervish et al. (2009), such a principle would rule out precisely the incoherent credence functions. However, as noted in Kelley and Neth 2023, it is not at all clear why only finite contexts are relevant to a contextual dominance principle. This seems like an ad hoc restriction. If the restriction is to be made to finite contexts, then there should be a good reason for doing so.

I suggest we quantify over contexts that admit of legitimate inaccuracy measures. Thus, we introduce the notion of measurability:

Definition 5. For some class of inaccuracy measures $\mathcal{I}$, let an opinion set $\mathcal{F}$ be $\mathcal{I}$-*measurable* if there is some inaccuracy measure in $\mathcal{I}$ defined on $\mathcal{F}$.

We do not assume that every opinion set is measurable relative to the "legitimate" class of inaccuracy measures, a class we denote by $\mathcal{L}$. But we do assume that a contextual dominance principle should only pay attention to accuracy dominance in $\mathcal{L}$-measurable contexts. This is motivated by the thought that if an opinion set does not admit of a legitimate inaccuracy measure, then how a credence function

[24] Thanks to Snow Zhang for helping me here.

does on that opinion set in terms of accuracy is not even well-defined and so could not be relevant to any principle of rationality in terms of accuracy. Note that it could turn out that only finite opinion sets admit of legitimate inaccuracy measures and so our contextual dominance principle in the end only pays attention to finite contexts (though see Proposition 8). But we would now have a principled reason for this focus, namely that the finite contexts are precisely those which admit of legitimate inaccuracy measures.

So we have a principled way of constraining which contexts will be quantified over in our contextual dominance principle. More precisely then, here is the proposed principle. Let c and d be credence functions on an opinion set $\mathcal{F}$. Say that c is *contextually dominated* by d if there is some $\mathcal{L}$-measurable context $\mathcal{S} \subseteq \mathcal{F}$ such that $c|_{\mathcal{S}}$ is strongly accuracy dominated by $d|_{\mathcal{S}}$ relative to some $\mathscr{I} \in \mathcal{L}$ defined on $\mathcal{S}$ and $d|_{\mathcal{S}'}$ is not weakly dominated by $c|_{\mathcal{S}'}$ relative to $\mathscr{I}$ for any context $\mathcal{S}'$ and $\mathscr{I} \in \mathcal{L}$ defined on $\mathcal{S}'$. In other words, by moving to d, one is guaranteed to be more accurate in some measurable context and not guaranteed to be less accurate in any measurable context. The *Contextual Dominance Principle* then states that one should not be contextually dominated by a credence function which is itself not contextually dominated.[25]

> **Contextual Dominance Principle** (CDP): If c is contextually dominated by d and d is not contextually dominated by any other credence function, then c is rationally impermissible.

As for the relationship between CDP and UD, consider a case where both apply, i.e., a case where c is defined on an $\mathcal{L}$-measurable opinion set. Then no matter the contents of $\mathcal{L}$, if c is ruled out by UD, then c is ruled out by CDP.[26] However, without further specification on $\mathcal{L}$, if c is ruled out by CDP, it does not follow that c is ruled out by UD. Thus, there is a sense in which CDP is a stronger principle than UD. There are, however, certain conditions on $\mathcal{L}$ for which CDP and UD are equivalent, e.g., if $\mathcal{L}$ contains precisely the inaccuracy measures meeting the conditions in Theorem 4. It is an interesting open question what conditions on $\mathcal{L}$ are necessary and sufficient for CDP to be equivalent to UD. This reduces to the question of what conditions on $\mathcal{L}$ are necessary and sufficient for the following to be true: if c (defined on a $\mathcal{L}$-measurable context) is dominated by an undominated credence function when restricted to some $\mathcal{L}$-measurable context, then c (unrestricted) is dominated by an undominated credence function.[27]

---

[25] We add the second clause for broadly the same reason that Pettigrew (2016a) argues it is *undominated* dominance which should be avoided: one's credence function is not obviously impermissible if it fails to meet a standard that no other credence function meets.

[26] Kelley and Neth (2023) worry that defending a local dominance principle like CDP will involve committing the fallacy of composition: an epistemic state can be irrational while having proper parts which, taken alone, are rationally permissible. My response to this worry is as follows: if a full opinion set is not $\mathcal{L}$-measurable, then we cannot apply UD since there's no legitimate notion of accuracy on the full opinion set; thus, all that is available to evaluation the credal state's rationality (from the perspective of accuracy) is the rationality of the credal state's proper parts. But if the full opinion set is $\mathcal{L}$-measurable, then CDP takes into account dominance relations on the full epistemic state, and so CDP will not mark as permissible a credence function which UD marks as impermissible. In short, the fallacy of composition objection only applies if the full opinion set is $\mathcal{L}$-measurable; and in this case violation of UD entails violation of CDP.

[27] Here is a partial result. Assume $\mathscr{I}(c, w) < \infty$ for any coherent credence function and world. Let c and d be defined on a countable $\mathcal{F}$. If $c|_{\mathcal{S}}$ is strongly dominated by $d|_{\mathcal{S}}$ for some $\mathcal{S} \subseteq \mathcal{F}$, then c is strongly dominated. Indeed, define d to be $d|_{\mathcal{S}}$ on $\mathcal{S}$ and c on $\mathcal{F} \setminus \mathcal{S}$. Then c is strongly dominated by d. There is no guarantee, however, that d is undominated or coherent.

Along with CDP, the contextual dominance argument assumes that there is an accuracy dominance argument that works in the finite setting. In particular, I will assume the following:

> **Finite admissibility** (FA): on every finite opinion set, there is a legitimate inaccuracy measure defined on it relative to which any incoherent credence function is strongly dominated by a credence function which is itself not even weakly dominated.[28]

Given that the task here is to offer a generalization of the accuracy dominance argument for Probabilism from the finite setting, it seems perfectly reasonable to assume that there is, in fact, an accuracy dominance argument for Probabilism in the finite setting. If there were not, it would make little sense to consider the infinite setting at all.

Finally, we turn to which necessary (but not necessarily sufficient) conditions we will assume of a legitimate inaccuracy measure. Ideally, these conditions would be as weak as possible so that the proposed argument is applicable for a wide range of answers to what is necessary and sufficient for legitimacy. I will offer results for two sets of properties. The first result, which will be proved in the Appendix, assumes only strict propriety.

Theorem 6. Assume finite admissibility and that every legitimate inaccuracy measure is strictly proper. Then a credence function on an arbitrary opinion set is contextually dominated if and only if it is incoherent.

The second result assumes a propriety condition and a continuity condition. For the propriety condition, I assume a stronger version of quasi-strict propriety, which I will call *total quasi-strict propriety*. While quasi-strict propriety requires that according to *some* extension of p, p is on average no more inaccurate than c for any c, total quasi-strict propriety requires that according to *all* extensions of p, p is on average no more inaccurate than c for any c. More formally, $\mathscr{I}$ satisfies total quasi-strict propriety if $\mathbb{E}_{p^*}\mathscr{I}(p,\cdot) \leq \mathbb{E}_{p^*}\mathscr{I}(c,\cdot)$ for all coherent credence functions p on $\mathcal{F}$, all probabilistic extensions p^* of p, and all credence functions c, with strict inequality if c is incoherent. The continuity condition that I assume can informally be described as extending the usual continuity condition to the infinite setting by requiring that when $\mathcal{F}$ is infinite, $\mathscr{I}(c,w)$ is continuous on $\mathcal{P}_{\mathcal{F}}$ for at least one of a broad range of reasonable topologies one might consider on $\mathcal{P}_{\mathcal{F}}$. Call this condition *permissive continuity*.[29] We then have the following result, which will be proved in the Appendix.[30]

---

[28]It follows from FA that on every finite opinion set, there is a legitimate inaccuracy measure defined on it relative to which any incoherent credence function is strongly dominated by a coherent credence function.

[29]More precisely, if $\mathcal{F}$ is finite, permissive continuity requires that for each $w \in W$, $\mathscr{I}(c,w)$ is continuous in the usual Euclidean topology on $\mathcal{C}_{\mathcal{F}}$. If $\mathcal{F}$ is infinite, consider the product topology, the sup-norm topology, and the ℓ^2 topology; the first two lead to natural continuity conditions on $\mathcal{C}_{\mathcal{F}}$ no matter the infinite opinion set; the third leads to a natural continuity condition in the countably infinite setting when coherent credences are square-summable. We say that $\mathscr{I}$ on $\mathcal{F}$ with $|\mathcal{F}| = \infty$ is permissively continuous if either i) for each $w \in W$, $\mathscr{I}(c,w)$ is continuous on $\mathcal{P}_{\mathcal{F}}$ assuming the product topology, ii) for each $w \in W$, $\mathscr{I}(c,w)$ is continuous on $\mathcal{P}_{\mathcal{F}}$ assuming the sup-norm topology, or iii) $\mathcal{F}$ is such that all coherent credence functions are square-summable and for each $w \in W$, $\mathscr{I}(c,w)$ is continuous on $\mathcal{P}_{\mathcal{F}}$ assuming the topology generated by the ℓ^2 norm. One might be able to establish the below result for an even weaker notion of continuity, but this will do for our purposes.

[30]Note that the conditions of Theorems 6 and 7 are of incomparable logical strength. A credence function can be totally quasi-strictly proper without being strictly proper; and a credence function can be strictly proper without being permissively continuous.

Theorem 7. Assume finite admissibility and that every legitimate inaccuracy measure is totally quasi-strictly proper and permissively continuous. Then a credence function on an arbitrary opinion set is contextually dominated if and only if it is incoherent.

So if we assume that $\mathcal{L}$ contains for each opinion set, precisely the additive, (permissively) continuous, and strictly proper inaccuracy measures—the standard conditions assumed in the finite case—then a credence function is contextually dominated if and only if it is incoherent. More generally, for any set of conditions stronger than either i) strict propriety or ii) total quasi-strict propriety and permissive continuity that one might conclude are necessary for a legitimate inaccuracy measure, Theorems 6 and 7 offer the basis for a fully general argument for Probabilism.

Putting all the pieces together, we have arrived at a Contextual Accuracy Dominance Argument for Probabilism.

<u>Contextual Accuracy Dominance Argument for Probabilism (CAP)</u>

Q1) Every legitimate inaccuracy measure satisfies either strict propriety or total quasi-strict propriety and permissive continuity.

Q2) Contextual Dominance Principle

Q3) Theorems 6 and 7

Q4) Finite admissibility

C) So, Probabilism.

CAP establishes Probabilism for any credence function on any opinion set and is thus a *genuinely* fully general accuracy dominance argument for Probabilism. Moreover, CAP has no premise assuming that inaccuracy can be measured on all opinion sets. Instead, we only assume that there is an accuracy dominance argument for Probabilism in the finite setting, an entirely uncontroversial assumption in the context of this discussion. This is very good news for advocates of Bayesian epistemology, Probabilism, and accuracy-first epistemology. Essentially no matter the strength of the necessary conditions on a legitimate inaccuracy measure, as long as they include minimally either strict propriety or total quasi-strict propriety and permissive continuity, CAP justifies Probabilism in full generality. And this is true regardless of whether every opinion set admits of a legitimate measure of inaccuracy.[31]

6. Further Directions

An important upshot of CAP is the "contextualizing strategy" that it exemplifies. Many existing accuracy arguments for various epistemic principles are restricted to the finite setting (e.g., Briggs

[31]The conditions that Nielsen assumes on a legitimate inaccuracy measure entail the conditions in Theorem 7; in this way, CAP can be seen as supplemental to and consistent with NGAP. Indeed, it is easy to see that weak continuity in expectation entails permissive continuity. As for total quasi-strict propriety, Nielsen does not discuss the question of whether the inequality in the propriety condition has to hold for all or only some extensions since he assumes credence functions are defined on power sets where there is only a single extension of a coherent credence function, namely itself. But we can, I think, interpret Nielsen as assuming total quasi-strict propriety without issue.

and Pettigrew 2020; Pettigrew 2016b). CAP shows that there exists a general strategy for extending these arguments to the infinite setting: assume an accuracy argument restricted to the finite setting uses accuracy-rationality bridge principle N requiring that one avoid some accuracy deficiency D; "contextualize" the argument by replacing N with the following contextualized version: one ought to avoid D *in $\mathcal{L}$-measurable contexts*. Adding in a version of finite admissability, one could potentially obtain a generalized version of the accuracy argument that applies even to infinite opinion sets, much in the way that CAP generalizes FGAP. It would be worth seeing how far this contextualizing strategy extends the arguments of accuracy-first epistemology.

One interesting mathematical question left open by everything said thus far is what the $\mathcal{L}$-measurable contexts are. This, of course, depends on the normative question of what the necessary and sufficient conditions are on a legitimate measure of inaccuracy. We do know that in the case that we take $\mathcal{L}$ to contain precisely the inaccuracy measures that satisfy strong additivity, permissive continuity, and strict propriety, some countably infinite opinion sets are $\mathcal{L}$-measurable:

Proposition 8. Let $\mathcal{F} = \{\{w\} : w \in W_0\} \subseteq 2^W$ with $|W \setminus W_0| = 1$ and W countably infinite. Then there exist strongly additive, permissively continuous, and strictly proper inaccuracy measures on $\mathcal{F}$.

The proof can be found in the Appendix.[32] The proof shows that, in fact, the natural extension of the *Brier score* to the countably infinite domain is one such inaccuracy measure.[33] The Brier score is a popular measure of inaccuracy in the finite setting, argued by some to be the unique legitimate measure of inaccuracy (Leitgeb and Pettigrew 2010a). There is more work to be done here, but Proposition 8 minimally shows that it is plausible that some infinite contexts are $\mathcal{L}$-measurable.

7. Conclusion

The project of giving a fully general accuracy dominance argument for Probabilism would greatly benefit from an in-depth investigation into the conditions required of a legitimate inaccuracy measure. As we have seen, in light of a number of impossibility results, attending to the infinite setting forces us to decide which conditions are truly non-negotiable. It is not clear at this point how this investigation will turn out. However, we should be very confident that we have a fully general accuracy-based argument for Probabilism regardless: so long as one of two sets of weak conditions come out as necessary conditions on a legitimate inaccuracy measure, we have available to us the Contextual Accuracy Dominance Argument for Probabilism.

Appendix

I begin by briefly reviewing a theory of finitely additive expectation for extended-real functions that are bounded below. I will follow the presentation in Schervish et al. 2020 and Nielsen 2023.

[32] Proposition 8 holds more generally for *compact* countably infinite opinion sets such that the σ-algebra generated by the opinion set is the power set and such that certain finiteness conditions are met. For further discussion and a proof of this more general result, see the Appendix.

[33] The Brier score of c on $\mathcal{F} \subseteq 2^W$ at $w \in W$ is $\sum_{p \in \mathcal{F}} (v_w(p) - c(p))^2$.

Given P a finitely additive probability function on 2^W, a finitely additive expectation with respect to P is defined on simple functions in the obvious way: for $\sum_{i=1}^n \alpha_i 1_{E_i}$ on W, set $\mathbb{E}_P(\sum_{i=1}^n \alpha_i 1_{E_i}) = \sum_{i=1}^n \alpha_i P(E_i)$. Since every bounded function can be uniformally approximated above and below by simple functions, for f a bounded function, we define

$$\mathbb{E}_P(f) = \sup\{\mathbb{E}_P(\varphi) : \varphi \leq f, \varphi \text{ simple}\} = \inf\{\mathbb{E}_P(\varphi) : f \leq \varphi, \varphi \text{ simple}\}.$$

For f potentially unbounded (though bounded below), we define

$$\mathbb{E}_P(f) = \sup\{\mathbb{E}_P(g) : g \leq f, g \text{ bounded}\}.$$

Here are two relevant features of finitely additive expectation as just defined (for a proof, see Lemma 9 of Schervish et al. 2020).

Proposition 9. Let f and g be extended-real functions bounded below on W, p, q finitely additive probability functions on 2^W, and $\lambda \in [0, 1]$. Then:

1. If $f \leq g$, then $\mathbb{E}_p(f) \leq \mathbb{E}_p(g)$;
2. $E_{\lambda p+(1-\lambda)q}(f) = \lambda\mathbb{E}_p(f) + (1-\lambda)\mathbb{E}_q(f)$.

We now prove the theorems that support the Contextual Accuracy Dominance Argument for Probabilism. First, note the following lemma, which can be proved from a number of existing results.

Lemma 10. If c is incoherent on $\mathcal{F}$, then c is incoherent on some $\mathcal{X} \subseteq \mathcal{F}$ with $|\mathcal{X}| < \infty$.

Proof. We prove the contrapositive: If c is coherent when restricted to every finite $\mathcal{X} \subseteq \mathcal{F}$, then c is coherent. First, note that c being coherent when restricted to every finite $\mathcal{X}$ entails that c is what Rao and Rao (1983) call a *positive real partial charge*. This follows from the fact that every finitely additive probability function on an algebra, when restricted to a subset of its domain, is a positive real partial charge (Proposition 3.2.3(b) of Rao and Rao 1983). According to Theorem 3.2.10 of Rao and Rao, c can thus be extended to a finitely additive probability function on $\mathcal{A}(\mathcal{F})$ and so is coherent. □

Theorem 6. Assume finite admissibility and that every legitimate inaccuracy measure is strictly proper. Then a credence function on an arbitrary opinion set is contextually dominated if and only if it is incoherent.

Proof. It is easy to see that (*) for any coherent c on $\mathcal{F}$, $\mathcal{L}$-measurable context $\mathcal{S} \subseteq \mathcal{F}$ and $\mathscr{I} \in \mathcal{L}$ defined on $\mathcal{S}$, $c|_\mathcal{S}$ is not weakly dominated relative to $\mathscr{I}$. This follows from strict propriety: if $c|_\mathcal{S}$ were weakly dominated by some $d|_\mathcal{S}$, then by Proposition 9.1, we have $E_{c|^*_\mathcal{S}}\mathscr{I}(d|_\mathcal{S}, \cdot) \leq \mathbb{E}_{c|^*_\mathcal{S}}\mathscr{I}(c|_\mathcal{S}, \cdot)$ for any extension $c|^*_\mathcal{S}$ of $c|_\mathcal{S}$. This contradicts that $\mathscr{I}$ is strictly proper.

Now assume that c is incoherent. Then by Lemma 10, c is incoherent on a finite subset $\mathcal{S} \subseteq \mathcal{F}$. By finite admissibility, $c|_\mathcal{S}$ is strongly dominated relative to a legitimate inaccuracy measure by some coherent credence function $d|_\mathcal{S}$ on $\mathcal{S}$. We can extend $d|_\mathcal{S}$ to a coherent credence function d on $\mathcal{F}$ (see, e.g., Corollary 3.2.10 from Rao and Rao 1983). Moreover, by (*), d is not weakly dominated on any

$\mathcal{L}$-measurable context, as it is coherent. Thus, c is contextually dominated by d. Next, assume c is coherent. By (*), c is not strongly dominated on any $\mathcal{L}$-measurable context and so is not contextually dominated. □

Theorem 7. Assume finite admissibility and that every legitimate inaccuracy measure is totally quasi-strictly proper and permissively continuous. Then a credence function on an arbitrary opinion set is contextually dominated if and only if it is incoherent.

Proof. We begin by showing that (*) for any coherent c on $\mathcal{F}$, $\mathcal{L}$-measurable context $\mathcal{S} \subseteq \mathcal{F}$ and $\mathscr{I} \in \mathcal{L}$ defined on $\mathcal{S}$, $c|_{\mathcal{S}}$ is not weakly dominated relative to $\mathscr{I}$. The proof follows the proof of Theorem 2 in Nielsen 2023 with some adjustments.

Assume toward a contradiction that there is some d on $\mathcal{F}$, $\mathcal{L}$-measurable context $\mathcal{S}$, and $\mathscr{I} \in \mathcal{L}$ defined on $\mathcal{S}$ such that $c|_{\mathcal{S}}$ is weakly dominated by some $d|_{\mathcal{S}}$. For ease of exposition, let p and q refer to $c|_{\mathcal{S}}$ and $d|_{\mathcal{S}}$ respectively. By weak dominance and Proposition 9.1, we have

$$\mathbb{E}_{\overline{p}}\mathscr{I}(q,\cdot) \leq \mathbb{E}_{\overline{p}}\mathscr{I}(p,\cdot) \tag{1}$$

for any probabilistic extension $\overline{p}$ of p. By weak dominance we also have for some w,

$$\mathscr{I}(q,w) < \mathscr{I}(p,w). \tag{2}$$

For $n \in \mathbb{N}$, let $p_n = n^{-1}v_w + (1-n^{-1})p$ be a coherent credence function on $\mathcal{S}$ and $\overline{p_n}$ be an extension of p_n to 2^W. Note that $\overline{p_n} = n^{-1}v_w + (1-n^{-1})\overline{p}$ for a probabilistic extension $\overline{p}$ of p. Indeed, for any n, $\overline{p} = \frac{\overline{p_n} - n^{-1}v_w}{1-n^{-1}}$ is a finitely additive probability function on 2^W extending p. By (1), Proposition 9.2, and total quasi-strict propriety we have for any n:

$$\begin{aligned}
n^{-1}\mathscr{I}(q,w) + (1-n^{-1})\mathbb{E}_{\overline{p}}\mathscr{I}(p) &\geq n^{-1}\mathscr{I}(q,w) + (1-n^{-1})\mathbb{E}_{\overline{p}}\mathscr{I}(q,\cdot) \\
&= \mathbb{E}_{\overline{p_n}}\mathscr{I}(q,\cdot) \\
&\geq \mathbb{E}_{\overline{p_n}}\mathscr{I}(p_n,\cdot) \\
&= n^{-1}\mathscr{I}(p_n,w) + (1-n^{-1})\mathbb{E}_{\overline{p}}\mathscr{I}(p_n,\cdot) \\
&\geq n^{-1}\mathscr{I}(p_n,w) + (1-n^{-1})\mathbb{E}_{\overline{p}}\mathscr{I}(p,\cdot)
\end{aligned}$$

By total quasi-strict propriety, $\mathbb{E}_{\overline{p}}\mathscr{I}(p) < \infty$, and so we can cancel out the second summands; it follows that $\mathscr{I}(q,w) \geq \mathscr{I}(p_n,w)$ for all n.

Now we show that $p_n \to p$ in the topologies referenced in permissive continuity. If $|\mathcal{F}| < \infty$ and so $\mathscr{I}$ is continuous assuming the Euclidean topology on $\mathcal{C}_{\mathcal{F}}$, then note that $p_n \to p$ in the Euclidean topology. If $|\mathcal{F}| = \infty$, then note that $p_n \to p$ in the product topology and the topology generated by the sup-norm. The former follows from the well-known fact that convergence in the product topology is equivalent to convergence of all of the projections of the sequence (Munkres 1975, Lemma 43.3).

As for the sup-norm topology, note that

$$||p_n - p||_\infty = ||n^{-1}v_w - n^{-1}p||_\infty = n^{-1}||v_w - p||_\infty,$$

which tends to 0 as $n \to \infty$, since $||v_w - p||_\infty \leq 1$. If $|\mathcal{F}| = \infty$ but countable and every coherent credence function on $\mathcal{F}$ is square-summable, then we must consider the case where $\mathscr{I}$ is continuous with respect to the topology generated by the ℓ^2 norm. Note then that $p_n \to p$ in the ℓ^2 norm. Indeed,

$$||p_n - p||_2 = ||n^{-1}v_w - n^{-1}p||_2 = n^{-1}||v_w - p||_2 \leq n^{-1}(||v_w||_2 + ||p||_2),$$

which tends to 0 as $n \to \infty$. Thus, no matter the topology relative to which $\mathscr{I}$ is continuous, we have $p_n \to p$ in that topology. Thus, since $\mathscr{I}(q, w) \geq \mathscr{I}(p_n, w)$ for all n and $p_n \to p$ in the required sense, by permissive continuity we have that $\mathscr{I}(q, w) \geq \mathscr{I}(p, w)$, which contradicts (2).

We now prove the theorem. First, assume that c is incoherent. Then by Lemma 10, c is incoherent on a finite subset $\mathcal{S} \subseteq \mathcal{F}$. By finite admissibility, $c|_\mathcal{S}$ is strongly dominated by some coherent credence function $d|_\mathcal{S}$ on $\mathcal{S}$. We can extend $d|_\mathcal{S}$ to a coherent credence function d on $\mathcal{F}$ (see, e.g., Corollary 3.2.10 from Rao and Rao 1983). Moreover, by (*), d is not weakly dominated on any $\mathcal{L}$-measurable context, as it is coherent. Thus, c is contextually dominated by d. Next, assume c is coherent. By (*), c is not strongly dominated on any $\mathcal{L}$-measurable context and so is not contextually dominated. □

To prove Proposition 8, we need to review material about *compact opinion spaces*. For a fuller discussion of compactness in this context, see Kelley 2023, Sec. 4.2.

Definition 11. An *opinion space* is a pair $(W, \mathcal{F})$, where W is a nonempty set and $\mathcal{F} \subseteq 2^W$.

Definition 12. Let $(W, \mathcal{F})$ be an opinion space. Let $f(n) \in \{0, 1\}$ and set $p_n^{f(n)} = p_n$ if $f(n) = 0$ and $p_n^{f(n)} = p_n^c$ if $f(n) = 1$. Then $(W, \mathcal{F})$ is *compact* if for any choice of $\{p_n\}_{n=1}^\infty \subseteq \mathcal{F}$ and $f : \mathbb{N} \to \{0, 1\}$, if $\bigcap_{n=1}^N p_n^{f(n)}$ is nonempty for every N, then $\bigcap_{n=1}^\infty p_n^{f(n)}$ is nonempty.

The key result about compact opinion spaces is that on compact opinion spaces, coherence is equivalent to *countable coherence*.

Definition 13. Let a credence function c be *countably coherent* if c extends to a countably additive probability function on a $\sigma(\mathcal{F})$. That is, there is a $c^* : \sigma(\mathcal{F}) \to [0, 1]$ such that:

1. $c^*(p) = c(p)$ for all $p \in \mathcal{F}$;
2. $c^*(\bigcup_{i=1}^\infty p_i) = \sum_{i=1}^\infty c^*(p_i)$ for $\{p_i\}_{i=1}^\infty \subseteq \mathcal{F}^*$ with $p_i \cap p_j = \varnothing$ for $i \neq j$;
3. $c^*(W) = 1$.

Otherwise, a credence function is *countably incoherent*.

Theorem 14 (Borkar et al. 2003)**.** The following are equivalent:

1. $(W, \mathcal{F})$ is compact;

2. for every credence function c on $(W, \mathcal{F})$, c is coherent if and only if c is countably coherent.

We prove a result more general than Proposition 8 first.

Proposition 15. Let $\mathcal{F}$ be countably infinite, compact, and such that $\sigma(\mathcal{F}) = 2^W$.[34] Assume also that each $c \in \mathcal{P}_{\mathcal{F}}$ is square-summable with $||c - v_w||_2$ bounded as a function of w. Then there exist strongly additive, permissively continuous, and strictly proper inaccuracy measures on $\mathcal{F}$.

Proof. Let $\mathcal{F} = \{p_i\}_{i=1}^{\infty}$ be as in the statement of the theorem. Consider $\mathscr{I}(c, w) = \sum_{i=1}^{\infty}(v_w(p_i) - c(p_i))^2$, the extension of the Brier score to the countably infinite setting. Clearly, $\mathscr{I}$ is strongly additive.

Turn now to permissive continuity. Note that since any coherent credence function on $\mathcal{F}$ is square-summable, we can consider the topology on $\mathcal{P}_{\mathcal{F}}$ induced by the ℓ^2 norm. Assume that $c_n \to c$ in the ℓ^2 norm. We need to show that $\mathscr{I}(c_n, w) \to \mathscr{I}(c, w)$ for any w.

Fix $w \in W$. By the triangle inequality, for any n:

$$||v_w - c_n||_2 \leq ||v_w - c||_2 + ||c - c_n||_2.$$

Squaring both sides, we have

$$\mathscr{I}(c_n, w) \leq \mathscr{I}(c, w) + ||c - c_n||_2^2 + 2||v_w - c||_2||c_n - c||_2.$$

Since $||c - c_n||_2^2 \to 0$ as $n \to \infty$ and $||v_w - c||_2 < \infty$, we have that $\lim_{n\to\infty} \mathscr{I}(c_n, w) \leq \mathscr{I}(c, w)$. For the other inequality, use the reverse triangle inequality to get:

$$||v_w - c_n||_2 \geq ||v_w - c||_2 - ||c_n - c||_2.$$

Squaring both sides, we have

$$\mathscr{I}(c_n, w) \geq \mathscr{I}(c, w) + ||c_n - c||_2^2 + 2||v_w - c||_2||c_n - c||_2.$$

Again, since $||c - c_n||_2^2 \to 0$ as $n \to \infty$ and $||v_w - c||_2 < \infty$, we have that $\lim_{n\to\infty} \mathscr{I}(c_n, w) \geq \mathscr{I}(c, w)$, establishing permissive continuity.

As for strict propriety, fix $c \in \mathcal{P}_{\mathcal{F}}$ and $d \in \mathcal{C}_{\mathcal{F}}$ with $d \neq c$. First, note that since $\mathcal{F}$ is compact, coherence is equivalent to countable coherence by Theorem 14. Thus, c can be extended to a countably additive probability function c^* on $\sigma(\mathcal{F}) = 2^W$. We need to show that $\mathbb{E}_{c^*}(c, \cdot) < \mathbb{E}_{c^*}(d, \cdot)$.

Since $||v_w - c||_2$ is a bounded function of W, by Proposition 9.1, $\mathbb{E}_{c^*}\mathscr{I}(c, \cdot) < \infty$. If $\mathbb{E}_{c^*}\mathscr{I}(d, \cdot) = \infty$, then clearly $\mathbb{E}_{c^*}\mathscr{I}(c, \cdot) < \mathbb{E}_{c^*}\mathscr{I}(d, \cdot)$. Consider next the case where $\mathbb{E}_{c^*}\mathscr{I}(d, \cdot) < \infty$.

It is well known that the Brier score for finite opinion sets is strictly proper. Thus, when $c(p_i) \neq$

[34] The third condition could be dropped if we defined propriety conditions to refer to extensions to arbitrary algebras. This would complicate the theory of finitely additive expectation, however.

$d(p_i)$:

$$\mathbb{E}_{c^*}(v_w(p_i) - c(p_i))^2 = c(p_i)(1-c(p_i))^2 + (1-c(p_i))(0-c(p_i))^2$$
$$< c(p_i)(1-d(p_i))^2 + (1-c(p_i))(0-d(p_i))^2 = \mathbb{E}_{c^*}(v_w(p_i) - d(p_i))^2.$$

If p_i is such that $d(p_i) = c(p_i)$, then clearly $\mathbb{E}_c(v_w(p_i) - c(p_i))^2 = \mathbb{E}_c(v_w(p_i) - d(p_i))^2$. Since $c \neq d$, there is at least one i such that $c(p_i) \neq d(p_i)$. Thus, using that $\mathbb{E}_{c^*}\mathscr{I}(c,\cdot) < \infty$ and $\mathbb{E}_{c^*}\mathscr{I}(d,\cdot) < \infty$, we have that:

$$\mathbb{E}_{c^*}\mathscr{I}(c,\cdot) = \mathbb{E}_{c^*}\sum_{i=1}^{\infty}(v_w(p_i) - c(p_i))^2 = \sum_{i=1}^{\infty}\mathbb{E}_{c^*}(v_w(p_i) - c(p_i))^2$$
$$< \sum_{i=1}^{\infty}\mathbb{E}_{c^*}(v_w(p_i) - d(p_i)) = \mathbb{E}_{c^*}\sum_{i=1}^{\infty}(v_w(p_i) - d(p_i)) = \mathbb{E}_{c^*}\mathscr{I}(d,\cdot),$$

establishing strict propriety. □

As a corollary, we have Proposition 8.

Proposition 8. Let $\mathcal{F} = \{\{w\} : w \in W_0\} \subseteq 2^W$ with $|W \setminus W_0| = 1$ and W countably infinite. Then there exist strongly additive, permissively continuous, and strictly proper inaccuracy measures on $\mathcal{F}$.

Proof. Note that $\mathcal{F}$ is compact. Indeed, there is only one infinite collection from $\{p : p \in \mathcal{F}\} \cup \{p^c : p \in \mathcal{F}\}$ with all non-empty finite intersections and such that there are infinitely many different elements in the collection: the collection $\{p^c : p \in \mathcal{F}\}$. But note that there is some $w \in \bigcap_{p \in \mathcal{F}} p^c$. It is also easy to see that $\sigma(\mathcal{F}) = 2^W$.

As for the finiteness conditions, note that for c to be coherent, it must be that $\sum_{w \in W_0} c(\{w\}) \leq 1$. Thus, c is square-summable. Moreover, if $w \in p_i$, we have $||v_w - c||_2 = \sqrt{(1-c(p_i))^2 + \sum_{i \neq j} c(p_j)^2}$; and if $w \notin p_i$ for all i, then $||v_w - c||_2 = \sqrt{\sum_{i=1}^{\infty} c(p_i)^2}$. Since c is square-summable, it follows that $||v_w - c||_2$ is bounded as a function of w. Thus, by Proposition 15, there exist strongly additive, permissively continuous, and strictly proper inaccuracy measures on $\mathcal{F}$. □

References

Vivek Borkar, Vijay Konda, and Sanjoy Mitter. On de Finetti coherence and Kolmogorov probability. *Statistics & Probability Letters*, 66:417–421, 2003.

R. A. Briggs and Richard Pettigrew. An accuracy-dominance argument for conditionalization. *Noûs*, 54(1):162–181, 2020. doi: 10.1111/nous.12258.

Catrin Campbell-Moore and Benjamin A Levinstein. Strict propriety is weak. *Analysis*, 81(1):8–13, 05 2021.

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.

Haim Gaifman and Marc Snir. Probabilities over rich languages, testing and randomness. *Journal of Symbolic Logic*, 47(3):495â548, 1982.

Sophie Horowitz. Accuracy and educated guesses. *Oxford Studies in Epistemology*, 6, 2019.

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.

Mikayla Kelley. On accuracy and coherence with infinite opinion sets. *Philosophy of Science*, 90(1): 92–128, 2023. doi: 10.1017/psa.2021.37.

Mikayla Kelley and Sven Neth. Accuracy and infinity: A dilemma for subjective Bayesians. *Synthese*, 201(12), 2023.

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

Hannes Leitgeb and Richard Pettigrew. An objective justification of Bayesianism II: The consequences of minimizing inaccuracy. *Philosophy of Science*, 77(2):236–272, 2010b. doi: 10.1086/651318.

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.

J.R. Munkres. *Topology: A First Course*. Prentice-Hall Inc., 1975.

Michael Nielsen. Accuracy and probabilism in infinite domains. *Mind*, 132(526):402–427, 2023. doi: 10.1093/mind/fzac053.

Ilho Park and Jaemin Jung. Infinite opinion sets and relative accuracy. *Journal of Philosophy*, 120 (6):285–313, 2023.

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*, 200(444), 2022. 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.

M. Schervish and T. Seidenfeld. An approach to consensus and certainty with increasing evidence. *Journal of Statistical Planning and Inference*, 25(3):401–414, 1990.

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.

Mark Schervish, Teddy Seidenfeld, Rafael Stern, and Joseph Kadane. What finite-additivity can add to decision theory. *Statistical Methods Applications*, 29:237–263, 2020.

Snow Zhang. Accuracy undefined. ms.