

ARTICLE

Accuracy-First Epistemology Without Additivity

Richard Pettigrew

Department of Philosophy, University of Bristol, Bristol, UK
Email: Richard.Pettigrew@bristol.ac.uk

(Received 22 August 2020; revised 16 November 2020; accepted 15 January 2021)

Abstract

Accuracy arguments for the core tenets of Bayesian epistemology differ mainly in the conditions they place on the legitimate ways of measuring the inaccuracy of our credences. The best existing arguments rely on three conditions: Continuity, additivity, and strict propriety. In this paper, I show how to strengthen the arguments based on these conditions by showing that the central mathematical theorem on which each depends goes through without assuming additivity.

1. Introduction

At the core of Bayesian epistemology lies a small number of fundamental credal principles. Probabilism tells you how your credences in logically related propositions should relate to one another. Conditionalization tells you how to update your credences in response to a specific sort of new evidence, namely, evidence that makes you certain of a proposition. The Principal Principle tells you how your credence about the objective chance of an event should relate to your credence in that event. And, for objective Bayesians, there are further principles that tell you how to set your prior credences, that is, those you have before you incorporate your data.

The accuracy-first programme in epistemology seeks new foundations for these central credal principles (Joyce 1998; Greaves and Wallace 2006; Pettigrew 2016a). The idea is straightforward. We adopt the orthodox Bayesian assumption that your uncertain doxastic states can be represented by assigning precise numerical credences to each proposition you consider. By convention, your credence in a proposition is represented by a single real number at least 0 and at most 1, which we take to measure how strongly you believe that proposition. We then represent the whole doxastic state by a credence function, which takes each proposition about which you have a credal opinion and assigns to it your credence in that proposition. So far, that's just the representational claim of orthodox Bayesianism. Now for the distinctive claim of accuracy-first epistemology. It is a claim about what makes a credal state, represented by a credence function, better or worse from the epistemic point of view. That is, it says what determines the epistemic or cognitive value of a credal state. It says: A credal state is better the more accurate it is; it is worse the more inaccurate it is. And we might intuitively think of the inaccuracy of a

credence function as how far it lies from the ideal credence function, which is the one that assigns maximal credence to each truth and minimal credence to each falsehood. In accuracy-first epistemology, we formulate mathematically precise ways of measuring this epistemic good. We then ask what principles a credence function should have if it is to serve the goal of attaining that good; or, perhaps better, what properties should it not have if it is to avoid being suboptimal from the point of view of that good. And indeed we can find arguments of exactly this sort for the various fundamental credal principles of Bayesianism that we listed above.

Each of these arguments has the same form. Its first premise specifies properties that an inaccuracy measure must have.¹ Its second premise provides a bridge principle connecting inaccuracy with rationality. And its third and final premise is a mathematical theorem that shows that, if we apply the bridge principle from the second premise using an inaccuracy measure that has the properties demanded by the first premise, it follows that any credence function that violates the credal principle we seek to establish is irrational. We thus conclude that principle.

So, for instance, Jim Joyce lays down a series of properties that legitimate measures of inaccuracy must have: Structure, extensionality, normality, dominance, weak convexity, and symmetry (Joyce 1998). He then formulates his bridge principle that connects inaccuracy and rationality: It says that a credence function is irrational if it is accuracy dominated; that is, if there is an alternative that is guaranteed to be more accurate. Then he proves a mathematical theorem to show that any credence function that is not probabilistic is accuracy dominated. And he concludes Probabilism.

Similarly, Hilary Greaves and David Wallace lay down a single property they take to be necessary for measure of inaccuracy (Greaves and Wallace 2006): It is strict propriety, and it will play a central role in what follows. Then they say that your updating plan is irrational if there is an alternative that your prior credence function expects to be more accurate. And finally, they prove that updating rules that proceed by conditioning your prior on your evidence, and only such rules, minimize expected inaccuracy from the point of view of your prior. And they conclude Conditionalization.

In contrast, R. A. Briggs and I demanded not only strict propriety, but also continuity and additivity. While Greaves and Wallace's bridge principle talks of the irrationality of your updating rule from the point of view of your prior, the principle that Briggs and I used talks of the irrationality of the combination of your prior and updating rule. It says that a prior and updating rule taken together are irrational if there is an alternative prior and an alternative updating rule that, when taken together, are more accurate than your prior and your updating rule taken together. And they show that any combination in which the updating rule does not proceed by applying Bayes' rule to the prior is rendered irrational by this bridge principle (Briggs and Pettigrew 2020).

My argument for the Principal Principle demanded the same three properties: Additivity, continuity, and strict propriety. The bridge principle: A credence function is irrational if there is an alternative that every possible objective chance function expects to be better. And the mathematical theorem says that there is such an alternative for every credence function that violates the Principal Principle (Pettigrew 2013).

¹ For various reasons, it's become standard in accuracy-first epistemology to work with measures of inaccuracy, rather than measures of accuracy. But there is no substantial difference: A measure of inaccuracy is simply the negative of a measure of accuracy, and vice versa.

Finally, my argument for the Principle of Indifference assumed two different properties: Egalitarianism and rendering indifference immodest. The bridge principle said that a credence function is irrational if there is an alternative whose worst-case inaccuracy is lower than the worst-case inaccuracy of the original one. And I showed that the only credence function that is not irrational by these lights is the uniform distribution, which is exactly what the Principle of Indifference demands (Pettigrew 2016b).

Now, an argument is only as strong as its premises are plausible. In this paper, I'd like to consider the first premise in each of these arguments. In these first premises, we lay down what we will demand of an inaccuracy measure. My aim is to take the current best version of this premise and improve it by making it less demanding. There are really eight sets of conditions offered in the literature:

- (I) In his 1998 argument for Probabilism, Joyce imposes six conditions on measures of inaccuracy: Structure, extensionality, normality, dominance, weak convexity, and symmetry (Joyce 1998).
- (II) In his 2009 argument for a restricted version of Probabilism, he imposes just two: Truth-directedness and coherent admissibility (Joyce 2009).
- (III) In their 2006 argument for Conditionalization, Greaves and Wallace impose just one: Strict propriety (Greaves and Wallace 2006).
- (IV) In their 2009 argument for Probabilism, Predd et al. impose three: Continuity, additivity, and strict propriety. These are also the three conditions imposed in my argument for the Principal Principle and the argument for Conditionalization that Briggs and I offer (Predd et al. 2009; Pettigrew 2013; Briggs and Pettigrew 2020).
- (V) In our 2010 arguments for Probabilism and Conditionalization, Leitgeb and I considered three different sets of conditions: For our purposes, the important condition is global normality and dominance, which entails additivity, the condition we seek to excise here (Leitgeb and Pettigrew 2010).
- (VI) In my 2014 argument for the Principle of Indifference, I imposed two conditions: The inaccuracy measures must be egalitarian and they must render indifference immodest (Pettigrew 2016b).
- (VII) In their 2016 argument, D'Agostino and Sinigaglia impose five: One-dimensional value sensitivity, subvector consistency, monotonic order sensitivity, permutation invariance, and replication invariance (D'Agostino and Dardanoni 2009; D'Agostino and Sinigaglia 2010).
- (VIII) In his 1982 paper, Lindley argued not for probabilism itself, but for a weaker principle. He did not assume, as we have, that credences are measured on a scale from 0 to 1, nor that 0 is the minimum and 1 is the maximum. Instead, he made few assumptions about the scale on which credences are measured; he imposed some reasonably weak conditions on measures of the inaccuracy of credences and then showed that those credal assignments that are not accuracy-dominated are precisely those that can be transformed into probabilistic credence functions using a particular transformation. However, while Lindley's conditions are weaker than some of the others listed here, they nonetheless include additivity (Lindley 1982).

There is a lot to be said about the relationships between these different sets of necessary conditions for inaccuracy measures, but that's not my purpose here. Here, I want to take what I think are the best accuracy-first arguments for Probabilism, Conditionalization, and the Principal Principle and improve them by weakening the demands they make of inaccuracy measures in their first premises. That is, I want to show that those arguments go through for a wider range of inaccuracy measures than we've previously allowed. As I will explain below, I take those best arguments to be the ones based on Predd et al.'s set of conditions: Strict propriety, continuity, and additivity. I will strengthen those arguments by showing that they go through if we impose only strict propriety and continuity. We do not need to impose the condition of additivity, which says roughly that the inaccuracy of a whole credence function should be the sum of the inaccuracies of the credences that it assigns. That is, we can strengthen those arguments by weakening their first premise.

Why should this interest us? After all, Joyce, as well as D'Agostino and Sinigaglia, have offered arguments for probabilism that don't assume additivity. True, but Patrick Maher (2002) has raised serious worries about Joyce's 1999 characterization, and I have built on those (Pettigrew 2016a, section 3.1); and Joyce's 2009 characterization applies only to credence functions defined over a partition and not those defined on a full algebra, so while its premises are weak, so is its conclusion. D'Agostino and Sinigaglia do not assume additivity, but their characterization does entail it, and the subvector consistency requirement they impose is implausibly strong for the same reason that additivity is implausibly strong. And, similarly, the global normality and dominance condition that Leitgeb and I assumed entails additivity and so is implausibly strong for the same reason. And, as noted above, Lindley explicitly assumes additivity. This suggests that the best existing accuracy-first argument for probabilism is the one based on Predd et al.'s results, which assumes additivity, strict propriety, and continuity. So there is reason to show that probabilism follows from strict propriety and continuity alone.

What about the Principal Principle? Well, the only existing accuracy-first argument for that is my 2013 argument, and that assumed additivity, strict propriety, and continuity. So, again, there is reason to show that the principle follows from strict propriety and continuity alone. What about Conditionalization? Here, Greaves and Wallace have offered an argument based on strict propriety alone—it does not assume additivity, nor even continuity. True, but their result applies to a very specific case, namely, one in which (i) you know ahead of time the partition from which your evidence will come, (ii) you know that you will learn a proposition iff it is true, and (iii) you form a plan for how you will respond should you learn a particular proposition from that partition. In contrast, the result that Briggs and I offered can be generalized to cover many more cases than just this. As we will see, it can be generalized to establish what I will call the Weak Reflection Principle, which entails the restricted version of Conditionalization that Greaves and Wallace consider. So, there is reason to excise additivity from the assumptions that Briggs and I made. However, as we will see, our argument assumes additivity in two guises: First, it demands that we measure the inaccuracy of an individual credence function using an inaccuracy measure that satisfies additivity; second, it assumes we obtain the inaccuracy of a combination of prior and updating rule by adding the inaccuracy of the prior to the inaccuracy of the

updating rule. We will see how to remove the first assumption of additivity, but not the second. That must await future work.

2. Predd et al.'s conditions

In this section, I describe Predd et al.'s set of conditions—the ones we numbered (IV) in our list above. This will furnish us with statements of strict propriety and continuity, the assumptions we'll use in our new arguments for Probabilism, the Principal Principle, and Conditionalization; and it will also introduce us to additivity, the assumption that we're dropping from the existing best arguments for these conclusions. We will explain the problems with additivity in section 3 below.

First, let's lay out the framework in which we're working:

- We write \mathcal{W} for the set of possible worlds. We assume \mathcal{W} is finite.² So $\mathcal{W} = \{w_1, \dots, w_n\}$.
- We write \mathcal{F} for the full algebra of propositions built over \mathcal{W} . That is, \mathcal{F} is the set of all subsets of \mathcal{W} .
- We write \mathcal{C} for the set of credence functions defined on \mathcal{F} . That is, \mathcal{C} is the set of functions $c : \mathcal{F} \rightarrow [0, 1]$.
- We write \mathcal{P} for the set of probabilistic credence functions defined on \mathcal{F} . That is, p is in \mathcal{P} iff p is in \mathcal{C} and (i) $p(\top) = 1$, and (ii) $p(A \vee B) = p(A) + p(B)$ when A and B are mutually exclusive, that is, when there is no possible world at which A and B are both true.
- Given a credence function c , we write c_i for the credence that c assigns to world w_i .
- We write w^i for the ideal credence function on \mathcal{F} at world w_i . That is, for X in \mathcal{F} ,

$$w^i(X) = \begin{cases} 1 & \text{if } X \text{ is true at } w_i \\ 0 & \text{if } X \text{ is false at } w_i \end{cases}$$

So, in particular, $w^i(w_j) = w_j^i = 1$ if $i = j$ and 0 if $i \neq j$.

- An inaccuracy measure is a function $\mathfrak{S} : \mathcal{C} \times \mathcal{W} \rightarrow [0, \infty]$. For c in \mathcal{C} and w_i in \mathcal{W} , $\mathfrak{S}(c, i)$ is the inaccuracy of c at world w_i .

Here are the three properties that Predd et al. demand of inaccuracy measures.

Continuity For each world w_i , $\mathfrak{S}(c, i)$ is a continuous function of c .

Additivity For each X in \mathcal{F} , there is a scoring rule

$\mathfrak{s}_X : \{0, 1\} \times [0, 1] \rightarrow [0, \infty]$ such that, for all c in \mathcal{C} and w_i in \mathcal{W} ,

$$\mathfrak{S}(c, i) = \sum_{X \in \mathcal{F}} \mathfrak{s}_X(w^i(X), c(X))$$

We say that the scoring rules \mathfrak{s}_X for each X in \mathcal{F} generate \mathfrak{S} .

Additivity says that the inaccuracy of a credence function is the sum of the inaccuracies of the individual credences it assigns.

Strict Propriety For all p in \mathcal{P} , $\sum_{i=1}^n p_i \mathfrak{S}(c, i)$ is minimized uniquely, as a function of c , at $c = p$.

² For consideration of the infinite case, see Kelley (2009).

That is, for all p in \mathcal{P} and $c \neq p$ in \mathcal{C} ,

$$\sum_{i=1}^n p_i \mathfrak{S}(p, i) < \sum_{i=1}^n p_i \mathfrak{S}(c, i)$$

Strict propriety says that each probabilistic credence function should expect itself to be most accurate.³

A few examples of inaccuracy measures:

- **Brier score** $\mathfrak{B}(c, i) = \sum_{X \in \mathcal{F}} (w^i(X) - c(X))^2$
- **Log score** $\mathfrak{L}(c, i) = -\log c_i$
- **Enhanced log score** $\mathfrak{L}^*(c, i) = \sum_{X \in \mathcal{F}} (-w^i(X) \log c(X) + c(X))$
- **Absolute value score** $\mathfrak{A}(c, i) = \sum_{X \in \mathcal{F}} |w^i(X) - c(X)|$
- **Logsumexp score**

$$\mathfrak{L} \mathfrak{S} \mathfrak{E}(c, i) = -\log(1 + \sum_{X \in \mathcal{F}} e^{c(X)}) + \frac{\sum_{X \in \mathcal{F}} (w^i(X) - c(X)) e^{c(X)}}{1 + \sum_{X \in \mathcal{F}} e^{c(X)}}$$

Then:

	Continuity	Additivity	Strict Propriety
\mathfrak{B}	✓	✓	✓
\mathfrak{L}	✓	×	×
\mathfrak{L}^*	✓	✓	✓
\mathfrak{A}	✓	✓	×
$\mathfrak{L} \mathfrak{S} \mathfrak{E}$	✓	×	✓

Some notes:

- The Brier score is additive. It is generated by using the quadratic scoring rule q for every proposition, where
 - $q(1, x) = (1 - x)^2$;
 - $q(0, x) = x^2$.
 Since q is continuous and strictly proper, so is \mathfrak{B} .
- The log score is not additive, and it is not strictly proper. A credence function that assigns credence 1 to each world dominates any credence function that assigns less than credence 1 to each world. The log score is, however, strictly \mathcal{P} -proper: That is, for all p in \mathcal{P} , $\sum_i p_i \mathfrak{L}(q, i)$ is minimized uniquely, among credence functions q in \mathcal{P} , at $q = p$.

³ A quick remark: We sometimes assume that the inaccuracy measure \mathfrak{S} satisfies additivity and then assume that the individual scoring rules \mathfrak{s}_X that generate \mathfrak{S} satisfy continuity and strict propriety, rather than assuming that \mathfrak{S} itself satisfies those conditions. The following fact shows that this makes no difference. We say that \mathfrak{s}_X is continuous if $\mathfrak{s}_X(1, x)$ and $\mathfrak{s}_X(0, x)$ are continuous functions of x . We say that \mathfrak{s}_X is strictly proper if, for any $0 \leq p \leq 1$, $p\mathfrak{s}_X(1, x) + (1 - p)\mathfrak{s}_X(0, x)$ is minimized, as a function of x , at $x = p$. Now, suppose \mathfrak{S} is additive with $\mathfrak{S}(c, i) = \sum_{X \in \mathcal{F}} \mathfrak{s}_X(w^i(X), c(X))$. Then
 (i) \mathfrak{S} is continuous iff each \mathfrak{s}_X is continuous.
 (ii) \mathfrak{S} is strictly proper iff each \mathfrak{s}_X is strictly proper.

- The enhanced log score is additive. It is generated using the enhanced logarithmic scoring rule \mathcal{I}^* for every proposition, where
 - $\mathcal{I}^*(1, x) = -\log x + x$;
 - $\mathcal{I}^*(0, x) = x$.
 Since \mathcal{I} is continuous and strictly proper, so is \mathcal{I}^* .
- The absolute value score is additive. It is generated by using the absolute scoring rule α for every proposition, where
 - $\alpha(1, x) = 1 - x$;
 - $\alpha(0, x) = x$.
 But α is not strictly proper. If $p < \frac{1}{2}$, then $p\alpha(1, x) + (1 - p)\alpha(0, x) = p(1 - x) + (1 - p)x$ is minimized at $x = 0$; if $p > \frac{1}{2}$, it is minimized at $x = 1$; if $p = \frac{1}{2}$, it is minimized at any $0 \leq x \leq 1$. So \mathcal{A} is not strictly proper either.
- The logsumexp score is strictly proper and continuous, but it is not additive. This shows that assuming strict propriety and continuity, as we do below, is strictly weaker than assuming strict propriety, additivity, and continuity. Lewis and Fallis (2019) give another example of an inaccuracy measure that satisfies strict propriety and continuity, but not additivity: They call it the asymmetric spherical rule.

My concern in this paper is to strengthen the accuracy-first arguments for Probabilism, the Principal Principle, and Conditionalization based on continuity + additivity + strict propriety by showing that they go through if we assume only continuity + strict propriety. Thus, for instance, they go through for the logsumexp score defined above and Lewis and Fallis' asymmetric spherical rule as well as the Brier and enhanced log scores. Of course, weakening the premises of an argument always strengthens it. So there seems good reason to note this fact regardless of your view of additivity. Nonetheless, in the next section, I explain why you might be suspicious of additivity. Then, in section 4, I give the arguments for Probabilism, the Principal Principle, and Conditionalization without appealing to it. In section 5, I conclude. The appendix contains the proofs of all the theorems on which these arguments are based.

3. Why additivity?

I should begin by pointing out that, while I appealed to Predd et al.'s mathematical results in my presentation of the accuracy dominance argument for probabilism in *Accuracy and the Laws of Credence*, those authors weren't themselves working in the accuracy-first framework (Pettigrew 2016a). What we are calling inaccuracy measures are for them loss functions; and in the context of loss functions, the additivity assumption is perfectly natural—providing the loss is stated in units of some commodity in which your utility is linear, the total loss to which your credence function is subject is the sum of the individual losses to which your credences are subject. But what about the accuracy-first framework? Is additivity still so natural there?⁴ Some claim it is not: Kotzen (2018, 778) claims that, by assuming additivity, we rule out certain plausible ways of measuring inaccuracy “more or less by fiat”; and,

⁴ Those already convinced that additivity is at least somewhat questionable might skip this section without loss.

as I mentioned above, Lewis and Fallis (2019) describe a continuous and strictly proper inaccuracy measure that seems reasonable, but note that it is not additive.

On the other hand, I wrote:

[S]umming the inaccuracy of individual credences to give the total inaccuracy of a credence function is the natural thing to do (Pettigrew 2016a, 49).

My reason? Your credence function is not a single, unified doxastic state, but rather merely the motley agglomeration of all of your individual credal doxastic states. We might mistakenly think of a credence function as unified because we represent it by a single mathematical object, but mathematical functions are anyway just collections of assignments of values to arguments. Or so I argued.

However, while this may be so, it does not entail that there are not features of a credal state that partly determine its accuracy but which cannot be captured by looking only at the credences individually—we might call these global features of the credence function. For instance, we might think that, at a world at which two propositions A and B have the same truth value, it is more accurate to have equal credence in A and in B than to have different credences in them. After all, the ideal credence function will have the same credence in them, namely, 1 in both or 0 in both. And perhaps resembling the ideal credence function in this respect gets you closer to it and therefore more accurate. But having the same credence in A and in B is a global feature of a credence function. To determine whether or not a credence function has it, you can't just look at the credences it assigns individually.

However, interestingly, while this is indeed a global feature of a credence function, some additive inaccuracy measures will in fact reward it. Take the Brier score, for instance. If A and B are both true, then the inaccuracy of assigning credences a and b to them, respectively, is $(1 - a)^2 + (1 - b)^2$. But it's easy to see that, if $a \neq b$, then assigning their average, $\frac{1}{2}a + \frac{1}{2}b$, to both is more accurate. Since $(1 - x)^2$ is a strictly convex function of x ,

$$(1 - a)^2 + (1 - b)^2 > \left(1 - \left(\frac{1}{2}a + \frac{1}{2}b\right)\right)^2 + \left(1 - \left(\frac{1}{2}a + \frac{1}{2}b\right)\right)^2$$

Similarly, if A and B are both false,

$$a^2 + b^2 > \left(\frac{1}{2}a + \frac{1}{2}b\right)^2 + \left(\frac{1}{2}a + \frac{1}{2}b\right)^2$$

So, if we measure inaccuracy using the Brier score, which is additive, having the global property in question—that is, assigning equal credences to A and B when A and B are either both true or both false—improves accuracy.

In *Accuracy and the Laws of Credence*, I argued that we should simply assume that, for any global feature of a credence function that we consider relevant to accuracy, it must be possible to capture it using additive measures along the lines just sketched.

Just as some hold that risk aversion phenomena in practical decision theory are best understood as the result of doing something other than maximizing expected utility—minimizing regret, for instance, or maximizing the quantity favoured by one of the many nonexpected utility theories—and not as having

a concave utility function, so any sensitivity to global features of credence functions ought to be understood either as following from their local features or as following from the adoption of an alternative decision principle and not as having a nonadditive inaccuracy measure. (Pettigrew 2016a, 51)

But why? Why think that this is true? I imagined that this follows from the fact that credence functions are not unified entities. I assumed that this warrants a defeasible assumption in favour of additivity. That assumption could be defeated if we were to find a global feature we wished to reward but which could not be rewarded by additive measures. But no such feature has presented itself.

Perhaps. But even if we grant the move from the disunified nature of credence functions to this defeasible assumption in favour of additivity, such a defeasible assumption is a flimsy basis for an argument, particularly since we have not systematically investigated the sorts of global features we might consider relevant to accuracy and so have rather sparse evidence that there is no defeater to be found. As we showed above, we can explain why one particular global feature is conducive to accuracy, namely, having the same credence in two propositions that have the same truth value. And indeed you can view the accuracy dominance argument for probabilism as furnishing us with another example. After all, probabilism makes two demands of a credence function. The first is local: The credence function should assign maximal credence to a tautology and minimal credence to a contradiction. The second is global: The credence it assigns to a disjunction of mutually exclusive propositions should be the sum of the credences it assigns to the disjuncts. To tell whether a credence function satisfies this latter condition, you must look at the relationships between the credences it assigns. However, the fact that you can run the accuracy dominance argument for probabilism using additive inaccuracy measures like the Brier score shows that you can show that this global feature of a credence function is conducive to accuracy without building into your inaccuracy measure that it should be rewarded explicitly. Indeed, that is one of the remarkable features of de Finetti's original proof, which is the ultimate source of Predd et al.'s theorem. But that's about it. Those are the only two global features of credence functions we've succeeded in capturing using additive inaccuracy measures.

So, in the end, to the extent that there still exists doubt that we have considered all global features of credence functions that are relevant to their accuracy and showed how their relevance can be captured by additive inaccuracy measures, there still exists doubt over additivity. And while there still exists doubt over additivity, removing that assumption from the arguments for Probabilism, the Principal Principle, and Conditionalization strengthens them.

4. Arguments without additivity

Let me begin this section by spelling out the arguments we wish to give. Then I'll move on to explaining and proving the theorems to which they appeal. As I noted above, each argument has the same form:

- (NC) Necessary conditions on being a legitimate inaccuracy measure.
- (BP) Accuracy-rationality bridge principle.

(MT) Mathematical theorem.

Therefore,

(PR) Principle of rationality.

The first component will be the same in each argument that we give. We will assume only that inaccuracy measures satisfy continuity and strict propriety. But the accuracy–rationality principles will differ from case to case. In the remainder of this section, I'll state each principle for which we're arguing and then the bridge principle to which we'll appeal in that argument.

To state some of the principles of credal rationality, we need some moderately technical notions. I'll lay them out here. Suppose \mathcal{X} is a set of credence functions. Then:

- \mathcal{X} is *convex* if it is closed under taking mixtures.
That is, for any c, c' in \mathcal{X} and any $0 \leq \lambda \leq 1$, $\lambda c + (1 - \lambda)c'$ is also in \mathcal{X} .
- \mathcal{X} is *closed* if it is closed under taking limits.
That is, for any infinite sequence c_1, c_2, \dots of credence functions in \mathcal{X} that tends to c in the limit, c is in \mathcal{X} .⁵
- The *convex hull* of \mathcal{X} is the smallest convex set that contains \mathcal{X} . We write it \mathcal{X}^+ .
That is, if \mathcal{Z} is convex and $\mathcal{X} \subseteq \mathcal{Z}$, then $\mathcal{X}^+ \subseteq \mathcal{Z}$.
- The *closed convex hull* of \mathcal{X} is the smallest closed set that contains the convex hull of \mathcal{X} . We write it $cl(\mathcal{X}^+)$.
Thus, if \mathcal{Z} is closed and $\mathcal{X}^+ \subseteq \mathcal{Z}$, then $cl(\mathcal{X}^+) \subseteq \mathcal{Z}$.

As well as clarifying the technical notions just presented, I also want to flag that these arguments appeal to a notion of possibility at various points: In setting up the framework, we have already talked of possible worlds and algebras built on those worlds; in accuracy dominance arguments for Probabilism, we will quantify over possible worlds; in the accuracy dominance argument for Conditionalization we quantify over possible worlds, but also over possible future credence functions that you currently endorse conditional on you receiving certain evidence; in expected inaccuracy arguments like Greaves and Wallace's argument for Conditionalization, we will sum over inaccuracies at possible worlds weighted by credences in those worlds to give expectations; and in chance dominance arguments for the Principal Principle, we'll quantify over possible objective chance functions. What notion of possibility is at play here? Joyce (1998) takes it to be logical possibility; in a recent paper (2020), I take it to be something like epistemic possibility. It is beyond the scope of this essay to argue for one or the other in detail, so I will leave the notion as a placeholder. But my own preference is for epistemic possibility. What makes it irrational to violate Probabilism, for instance, is that there is another credence function that is more accurate than yours at all epistemically possible worlds—that is, you can

⁵ We say that c_1, c_2, \dots tends to c in the limit—written $\lim_{n \rightarrow \infty} c_n = c$ —if

$$(\forall \varepsilon > 0)(\exists n)(\forall m > n)(\forall X \in \mathcal{F})[|c_m(X) - c(X)| < \varepsilon]$$

tell from the inside, so to speak, that this alternative is better than yours because the only worlds you need to consider are those you consider possible.

4.1 Argument for Probabilism

Here's probabilism:

Probabilism. Your credence function at each time during your epistemic life should be probabilistic.

That is, if c is your credence function at a given time, then we should have:

- (i) $c(\top) = 1$;
- (ii) $c(A \vee B) = c(A) + c(B)$ for mutually exclusive A and B .

And here's the bridge principle: A credence function is irrational if there's another that is guaranteed to be more accurate; that is, if it is accuracy dominated. That is:

Worldly dominance. If there is c^* such that, for all possible worlds w_i , $\mathfrak{S}(c^*, i) < \mathfrak{S}(c, i)$, then c is irrational.

Thus, to move from the claim that an inaccuracy measure must satisfy continuity and strict propriety, together with worldly dominance, to probabilism, we need the following theorem:

Theorem 1. *Suppose \mathfrak{S} is continuous and strictly proper. If c is not in \mathcal{P} , there is c^* in \mathcal{P} , such that $I(c^*, i) < \mathfrak{S}(c, i)$ for all w_i in \mathcal{W} .*

We prove that in the appendix.

4.2 Argument for the Principal Principle

The principal principle is usually stated as follows, where, for any probability function ch , C_{ch} is the proposition that says that ch is the objective chance function:⁶

Principal principle. If C_{ch} is in \mathcal{F} and $c(C_{ch}) > 0$, then, for all A in \mathcal{F} , we should have

$$c(A|C_{ch}) = ch(A)$$

The version we'll consider here is more general than this version. And indeed, the more general version entails the usual version.

⁶ This formulation follows Ned Hall's (Hall 2004). It differs somewhat from David Lewis' original formulation (Lewis 1980). For one thing, it does not require a notion of admissible evidence. Lewis' version requires that, if E is admissible for A , then we should have $c(A|chance\ of\ A\ is\ r\ \&\ E) = r$. Hall's version entails that, for any E , $c(A|C_{ch} \ \&\ E) = ch(A|E)$.

Generalized Principal Principle. Suppose \mathcal{A} is the set of possible objective chance functions. Then your credence function should be in the closed convex hull of \mathcal{A} .

Thus, for instance, if all the possible objective chance functions consider A and B equally likely, your credence function should consider them equally likely; and if all the possible objective chance functions consider A more likely than B , then you should consider A at least as likely as B ; and if all the possible objective chance functions consider A to be between 30% and 70% likely, then you should not assign credence 0.8 to A ; and so on.

Suppose ch is a possible objective chance function and C_{ch} is in \mathcal{F} . And suppose further that ch is not a self-undermining chance function. That is, it is certain that it gives the chances. That is, $ch(C_{ch}) = 1$. Then, if you satisfy the Generalized Principal Principle, then $c(A|C_{ch}) = ch(A)$.⁷

So that's the version of the Principal Principle that we'll consider. And here's the bridge principle: A credence function is irrational if there is another that is guaranteed to have greater expected accuracy from the point of view of the objective chance function. That is:

Chance Dominance. If there is c^* such that, for all possible chance functions ch , $\sum_i ch_i \mathfrak{S}(c^*, i) < \sum_i ch_i \mathfrak{S}(c, i)$, then c is irrational.

Thus, to move from the claim that an inaccuracy measure must satisfy continuity and strict propriety, together with Chance Dominance, to the Generalized Principal Principle, we need the following theorem:

Theorem 2. Suppose \mathfrak{S} is continuous and strictly proper. If c is not in $cl(\mathcal{A}^+)$, there is c^* in $cl(\mathcal{A}^+)$, such that $\sum_i ch_i \mathfrak{S}(c^*, i) < \sum_i ch_i \mathfrak{S}(c, i)$, for all ch in \mathcal{A} .

Again, we prove this in the appendix.

What's more, Theorem 2 allows us to improve not only my argument for the Principal Principle, but also my argument for linear pooling as the correct method of probabilistic judgment aggregation (Pettigrew 2019a). In that argument, I appealed to the following bridge principle: If we have a group of individuals with probabilistic credence functions c_1, \dots, c_m , and we wish to find a credence function that aggregates them, then we should not pick an aggregate c if there is an alternative c^* such that each member of the group expects c^* to be more accurate than they expect c to be. I assumed that inaccuracy measures satisfy additivity, continuity, and strict propriety. And I showed that, if you use an aggregation method other than linear pooling, then there will be an alternative aggregate that everyone expects to be better; but if you use linear pooling, there won't be. Armed with Theorem 2, we obtain the same result assuming only continuity and strict propriety and thereby strengthen my original argument.

⁷ To see this, note that having this property is closed under taking mixtures and taking limits.

4.3 Argument for Conditionalization and the Weak Reflection Principle

The version of conditionalization for which Greaves and Wallace as well as Briggs and I argue is this:

Plan Conditionalization Suppose your prior is c_0 and you know that your evidence will come from a partition E_1, \dots, E_m . And suppose you will learn a particular cell of this partition iff it is true. And suppose you plan to adopt credence function c_k in response to evidence E_k . Then, for all E_k with $c_0(E_k) > 0$ and for all X in \mathcal{F} , we should have

$$c_k(X) = c_0(X|E_k) = \frac{c_0(XE_k)}{c_0(E_k)}$$

That is, you should plan to update by applying Bayes' rule to your prior.

As we mentioned above, we won't argue for Plan Conditionalization directly. Rather, we'll argue for a more general principle of rationality called the Weak Reflection Principle; and we'll show that it entails Plan Conditionalization (Pettigrew 2020).

Weak Reflection Principle. Suppose $\mathcal{R} = \{c_1, \dots, c_m\}$ is the set of possible future credence functions you endorse. Then your current credence function should be in the convex hull of \mathcal{R} .

Here's the idea behind this principle. A lot might happen between now and tomorrow. You might see new sights, think new thoughts; you might forget things you know today, take mind-altering drugs that enhance or impair your thinking; and so on. So perhaps there is a set of credence functions, any one of which you think you might have tomorrow. Some of those you'll endorse—perhaps those you'd adopt if you saw certain new things, or enhanced your cognition in various ways. And some of them you'll disavow—perhaps those that you'd adopt if you were to forget certain things, or were to impair your cognition in some way. The Weak Reflection Principle asks you to separate out the wheat from the chaff, and once you've identified the ones you endorse, it tells you that your current credence function should lie within the convex hull of those future ones.

Now, suppose that you are in the situation that Plan Conditionalization covers. That is, (i) you know that you will receive evidence from the partition E_1, \dots, E_m , (ii) you will learn E_k iff E_k is true, and (iii) you form a plan for how to respond to these different possible pieces of evidence—you will adopt c_1 if you learn E_1 , c_2 if you learn E_2 , and so on. Thus, the possible future credence functions that you endorse are c_1, \dots, c_m , for they are the possible outcomes of a plan that you have made and which you know covers all the bases. Then, by the Weak Reflection Principle, c_0 should be in the convex hull of c_1, \dots, c_m . Then, if $c_k(E_k) = 1$ and $c_0(E_k) > 0$, then $c_k(X) = c_0(X|E_k)$, as Plan Conditionalization requires. After all, by the Weak Reflection Principle, there are $0 \leq \lambda_1, \dots, \lambda_m \leq 1$ with $\sum_{i=1}^m \lambda_i = 1$ such that $c_0(X) = \sum_{i=1}^m \lambda_i c_i(X)$ for all X in \mathcal{F} . What's more, by assumption, $c_i(E_k) = 1$ if $i = k$ and $c_i(E_k) = 0$ if $i \neq k$. So, $c_i(XE_k) = c_i(X)$ if $i = k$ and $c_i(XE_k) = 0$ if $i \neq k$. So,

$$c_0(X|E_k) = \frac{\sum_{i=1}^m \lambda_i c_i(XE_k)}{\sum_{i=1}^m \lambda_i c_i(E_k)} = \frac{\lambda_k c_k(XE_k)}{\lambda_k c_k(E_k)} = c_k(X)$$

as required.

However, the Weak Reflection Principle applies in many other situations beyond those imagined by Plan Conditionalization. For instance, suppose you know you'll receive evidence from E_1, \dots, E_m , but, while these propositions are mutually exclusive, they do not exhaust the logical space. Then again, if you plan to adopt c_k upon learning E_k , and $c_k(E_k) = 1$ and $c_0(E_k) > 0$, then the Weak Reflection Principle says that you should plan so that $c_k(X) = c_0(X|E_k)$. That follows from the working at the end of the previous paragraph. Or suppose your evidence is not perfectly reliable. So it's not certain that you'll learn E_k iff E_k is true. Then again, if you plan to adopt c_k upon learning E_k , and $c_k(E_k) = 1$ and $c_0(E_k) > 0$, then the Weak Reflection Principle says that you should plan so that $c_k(X) = c_0(X|E_k)$. Again, this follows from the working above. What's more, suppose you know you'll receive evidence from E_1, \dots, E_m , but you don't make a deterministic plan. Perhaps you plan as follows: If I learn E_1 , then I'll adopt c_1 or I'll adopt c'_1 ; if I learn E_2 , then I'll adopt c_2 or I'll adopt c'_2 ; and so on. Then Plan Conditionalization no longer applies. But the Weak Reflection Principle does. It constrains your plans. It says that your current credence function should be in the convex hull of $c_1, c'_1, c_2, c'_2, \dots, c_m, c'_m$ (Pettigrew 2019b).

So that's the Weak Reflection Principle. Now for the bridge principle we use to establish it: It is irrational for you to have a prior credence function and a set of possible future credence functions you endorse if there is some alternative prior and, for each possible future credence function you endorse, an alternative to that, such that the sum of the inaccuracy of your prior and the inaccuracy of one of the possible future credence functions you endorse is always greater than the inaccuracy of the alternative prior and the corresponding alternative possible future credence function. That is:

Diachronic Worldly Dominance. Suppose c_0 is your current credence function and c_1, \dots, c_m are the possible future credence functions you endorse. Then if there are $c_0^*, c_1^*, \dots, c_m^*$ such that, for each $1 \leq k \leq m$ and all w_i in \mathcal{W} ,

$$\mathfrak{I}(c_0^*, i) + \mathfrak{I}(c_k^*, i) < \mathfrak{I}(c_0, i) + \mathfrak{I}(c_k, i)$$

then you are irrational.

Thus, to move from the claim that an inaccuracy measure must satisfy continuity and strict propriety, together with diachronic worldly dominance, to the Weak Reflection Principle, we need the following theorem:

Theorem 3. Suppose \mathfrak{I} is continuous and strictly proper. If c_0 is not in $\{c_1, \dots, c_n\}^+$, then there are $c_0^*, c_1^*, \dots, c_n^*$ such that, for all w_i in \mathcal{W} ,

$$\mathfrak{I}(c_0^*, i) + \mathfrak{I}(c_k^*, i) < \mathfrak{I}(c_0, i) + \mathfrak{I}(c_k, i)$$

We prove this in the appendix. Note, however, as we mentioned above: While our new proof removes one way in which an additivity enters into the argument that Briggs and I gave, it doesn't remove another. We still take the inaccuracy of your prior and your posterior, taken together, to be the sum of the inaccuracy of your prior and the

inaccuracy of your posterior. It must await future work to explore whether that assumption might also be removed.

5. Conclusion

Accuracy arguments for the core Bayesian tenets differ mainly in the conditions they place on the legitimate inaccuracy measures. The best existing arguments rely on Predd et al.'s conditions: Continuity, additivity, and strict propriety. In this paper, I showed how to strengthen each argument based on these by showing that the central mathematical theorem on which it depends goes through without assuming additivity.

Acknowledgments. I am very grateful to Catrin Campbell-Moore and Jason Konek for very valuable discussions around this topic and to the two referees for this journal for extremely useful feedback.

References

- Briggs, R. A., and Richard Pettigrew. 2020. "An Accuracy-Dominance Argument for Conditionalization." *Noûs* 54 (1):162–81.
- D'Agostino, Marcello, and Valentino Dardanoni. 2009. "What's So Special about Euclidean Distance? A Characterization with Applications to Mobility and Spatial Voting." *Social Choice and Welfare* 33 (2):211–33.
- D'Agostino, Marcello, and Corrado Sinigaglia. 2010. "Epistemic Accuracy and Subjective Probability." In *EPSA Epistemology and Methodology of Science: Launch of the European Philosophy of Science Association*, ed. Mauricio Suárez, Mauro Dorato, and Miklós Rédei, 95–105. Netherlands: Springer.
- Greaves, Hilary, and David Wallace. 2006. "Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility." *Mind* 115 (459):607–32.
- Hall, Ned. 2004. "Two Mistakes about Credence and Chance." *Australasian Journal of Philosophy* 82 (1):93–111.
- Joyce, James M. 1998. "A Nonpragmatic Vindication of Probabilism." *Philosophy of Science* 65 (4):575–603.
- Joyce, James M. 2009. "Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief." In *Degrees of Belief*, ed. Franz Huber and Christoph Schmidt-Petri. Springer.
- Kelley, Mikayla (ms). "On Accuracy and Coherence with Infinite Opinion Sets." Unpublished manuscript.
- Kotzen, Matthew. 2018. "Comments on Richard Pettigrew's *Accuracy and the Laws of Credence*." *Philosophy and Phenomenological Research* 96 (3):776–83.
- Leitgeb, Hannes, and Richard Pettigrew. 2010. "An Objective Justification of Bayesianism I: Measuring Inaccuracy." *Philosophy of Science* 77:201–35.
- Lewis, David. 1980. "A Subjectivist's Guide to Objective Chance." In *Studies in Inductive Logic and Probability*, vol. II, ed. R. C. Jeffrey. Berkeley: University of California Press.
- Lewis, Peter J., and Don Fallis. 2019. "Accuracy, Conditionalization, and Probabilism." *Synthese* 198 (5):4017–33.
- Lindley, Dennis V. 1982. "Scoring Rules and the Inevitability of Probability." *International Statistical Review* 50:1–26.
- Maher, Patrick. 2002. "Joyce's Argument for Probabilism." *Philosophy of Science* 69 (1):73–81.
- Pettigrew, Richard. 2013. "A New Epistemic Utility Argument for the Principal Principle." *Episteme* 10 (1):19–35.
- Pettigrew, Richard. 2016a. *Accuracy and the Laws of Credence*. Oxford: Oxford University Press.
- Pettigrew, Richard. 2016b. "Accuracy, Risk, and the Principle of Indifference." *Philosophy and Phenomenological Research* 92 (1):35–59.
- Pettigrew, Richard. 2019a. "On the Accuracy of Group Credences." In *Oxford Studies in Epistemology*, vol. 6, ed. Tamar Szabo Gendler and John Hawthorne, 137–160. Oxford: Oxford University Press.
- Pettigrew, Richard. 2019b. "What Is Conditionalization, and Why Should We Do it?" *Philosophical Studies* 177 (11):3427–63.
- Pettigrew, Richard. 2020. "Logical Ignorance and Logical Learning." *Synthese* 198 (10):9991–10020.
- Pettigrew, Richard (ta). "Bayesian Updating When What You Learn Might Be False." *Erkenntnis*. <https://doi.org/10.1007/s10670-020-00356-8>.

Predd, Joel B., Robert Seiringer, Elliott H. Lieb, Daniel N. Osherson, H. Vincent Poor, and Sanjeev R. Kulkarni. 2009. "Probabilistic Coherence and Proper Scoring Rules." *IEEE Transactions of Information Theory* 55 (10):4786–92.

Rudin, Walter. 1976. *Principles of Mathematical Analysis*. International Series in Pure and Applied Mathematics. New York: McGraw-Hill.

Savage, Leonard J. 1971. Elicitation of Personal Probabilities and Expectations. *Journal of the American Statistical Association* 66 (336):783–801.

6. Appendix: The proofs

6.1 Theorems 1 and 2

As will be obvious to anyone familiar with the proof strategy in (Predd et al. 2009), many of the ideas used here are adapted from that paper, and those in turn are adapted from insights in (Savage 1971). Predd et al. proceed by proving a connection between additive and continuous strictly proper inaccuracy measures, on the one hand, and a sort of divergence between credence functions, on the other. A divergence from one credence function defined on \mathcal{F} to another is a function $\mathfrak{D} : \mathcal{C} \times \mathcal{C} \rightarrow [0, \infty]$ such that

- (i) if $c \neq c'$, then $\mathfrak{D}(c, c') > 0$,
- (ii) if $c = c'$, then $\mathfrak{D}(c, c') = 0$.

That is, the divergence from one credence function to another is always positive, while the divergence from a credence function to itself is always zero.

Given a continuous scoring rule \mathfrak{s}_X , Predd et al. define the following function for $0 \leq x \leq 1$:

$$\varphi_X(x) := -x\mathfrak{s}_X(1, x) - (1 - x)\mathfrak{s}_X(0, x)$$

And then, given an additive and continuous strictly proper inaccuracy measure \mathfrak{I} that is generated by continuous strictly proper scoring rules \mathfrak{s}_X for X in \mathcal{F} , they define a divergence as follows:

$$\mathfrak{D}(c, c') = \sum_{X \in \mathcal{F}} \varphi_X(c(X)) - \sum_{X \in \mathcal{F}} \varphi_X(c'(X)) - \sum_{X \in \mathcal{F}} \varphi'_X(c'(X))(c(X) - c'(X))$$

They show that this is a species of divergence known as a Bregman divergence. What's more, using a representation theorem due to Savage (1971), they show that, for any w_i in \mathcal{W} and c in \mathcal{C} , $\mathfrak{D}(w^i, c) = \mathfrak{I}(c, i)$. That is, the divergence from the ideal credence function at a world to a given credence function is the inaccuracy of the given credence function at that world. Having established this, Predd et al. can then appeal to various properties of Bregman divergences to establish their dominance result. In our proofs, since we do not assume additivity, it is not so straightforward to construct a Bregman divergence from our inaccuracy measures. Instead, we construct a restricted divergence, which gives the divergence from a probabilistic credence function to any credence function. We say that it is restricted because it doesn't say how to measure the divergence from a nonprobabilistic credence function to another credence function. That is, our divergence is defined on $\mathcal{P} \times \mathcal{C}$, not on $\mathcal{C} \times \mathcal{C}$. As a result, the function we define is not a Bregman divergence. And, as a result of that, we must prove for ourselves that it has the properties needed to establish the

various theorems that have previously been proved using the Bregman divergences that Predd et al. construct from additive inaccuracy measures. The following lemma defines this restricted divergence and describes four properties it boasts.

Lemma 4. *Suppose \mathfrak{S} is a strictly proper inaccuracy measure. Then define $\mathfrak{D}_{\mathfrak{S}} : \mathcal{P} \times \mathcal{C} \rightarrow [0, \infty]$ as follows:*

$$\mathfrak{D}_{\mathfrak{S}}(p, c) := \sum_i p_i \mathfrak{S}(c, i) - \sum_i p_i \mathfrak{S}(p, i)$$

Then:

- (i) $\mathfrak{D}_{\mathfrak{S}}$ is a divergence. That is, $\mathfrak{D}_{\mathfrak{S}}(p, c) \geq 0$ for all p in \mathcal{P} and c in \mathcal{C} with equality iff $p = c$.
- (ii) $\mathfrak{D}_{\mathfrak{S}}(w^i, c) = \mathfrak{S}(c, i)$, for all $1 \leq i \leq n$.
- (iii) $\mathfrak{D}_{\mathfrak{S}}$ is strictly convex in its first argument. That is, for all p, q in \mathcal{P} and c in \mathcal{C} and for all $0 \leq \lambda \leq 1$,

$$\mathfrak{D}_{\mathfrak{S}}(\lambda p + (1 - \lambda)q, c) < \lambda \mathfrak{D}_{\mathfrak{S}}(p, c) + (1 - \lambda)\mathfrak{D}_{\mathfrak{S}}(q, c)$$

- (iv) $\mathfrak{D}_{\mathfrak{S}}(p, c) \geq \mathfrak{D}_{\mathfrak{S}}(p, q) + \mathfrak{D}_{\mathfrak{S}}(q, c)$ iff $\sum_{i=1}^n (p_i - q_i)(\mathfrak{S}(c, i) - \mathfrak{S}(q, i)) \geq 0$

Proof of Lemma 4.

- (i) \mathfrak{S} is strictly proper. So:

- (a) if $c \neq p$, $\sum_i p_i \mathfrak{S}(p, i) < \sum_i p_i \mathfrak{S}(c, i)$; and,
- (b) if $c = p$, $\sum_i p_i \mathfrak{S}(p, i) = \sum_i p_i \mathfrak{S}(c, i)$.

So, $\mathfrak{D}_{\mathfrak{S}}(p, c) = \sum_i p_i \mathfrak{S}(c, i) - \sum_i p_i \mathfrak{S}(p, i) \geq 0$ with equality iff $c = p$.

- (ii)

$$\begin{aligned} \mathfrak{D}_{\mathfrak{S}}(w^i, c) = & 0 \times \mathfrak{S}(c, 1) + \dots + \\ & 0 \times \mathfrak{S}(c, i - 1) + 1 \times \mathfrak{S}(c, i) + 0 \times \mathfrak{S}(c, i + 1) + \dots + \\ & 0 \times \mathfrak{S}(c, n) - \\ & 0 \times \mathfrak{S}(w^i, 1) - \dots - \\ & 0 \times \mathfrak{S}(w^i, i - 1) - 1 \times \mathfrak{S}(w^i, i) - 0 \times \mathfrak{S}(w^i, i + 1) - \dots - \\ & 0 \times \mathfrak{S}(w^i, n) = \mathfrak{S}(c, i) \end{aligned}$$

since $\mathfrak{S}(w^i, i) = 0$.

- (iii) Suppose p and q are in \mathcal{P} , and suppose $0 < \lambda < 1$. Then let $r = \lambda p + (1 - \lambda)q$. Then, since $\sum_i p_i \mathfrak{S}(c, i)$ is uniquely minimized, as a function of c , at $c = p$, and $\sum_i q_i \mathfrak{S}(c, i)$ is uniquely minimized, as a function of c , at $c = q$, we have

$$\sum_i p_i \mathfrak{S}(p, i) < \sum_i p_i \mathfrak{S}(r, i)$$

$$\sum_i q_i \mathfrak{S}(q, i) < \sum_i q_i \mathfrak{S}(r, i)$$

Thus

$$\lambda[-\sum_i p_i \mathfrak{S}(p, i)] + (1 - \lambda)[- \sum_i q_i \mathfrak{S}(q, i)] >$$

$$\lambda[-\sum_i p_i \mathfrak{S}(r, i)] + (1 - \lambda)[- \sum_i q_i \mathfrak{S}(r, i)] =$$

$$-\sum_i r_i \mathfrak{S}(r, i)$$

Now, adding

$$\lambda \sum_i p_i \mathfrak{S}(c, i) + (1 - \lambda) \sum_i q_i \mathfrak{S}(c, i) =$$

$$\sum_i (\lambda p_i + (1 - \lambda) q_i) \mathfrak{S}(c, i) = \sum_i r_i \mathfrak{S}(c, i)$$

to both sides gives

$$\lambda[\sum_i p_i \mathfrak{S}(c, i) - \sum_i p_i \mathfrak{S}(p, i)] +$$

$$(1 - \lambda)[\sum_i q_i \mathfrak{S}(c, i) - \sum_i q_i \mathfrak{S}(q, i)] >$$

$$\sum_i r_i \mathfrak{S}(c, i) - \sum_i r_i \mathfrak{S}(r, i)$$

That is,

$$\lambda \mathfrak{D}_{\mathfrak{S}}(p, c) + (1 - \lambda) \mathfrak{D}_{\mathfrak{S}}(q, c) > \mathfrak{D}_{\mathfrak{S}}(r, c) = \mathfrak{D}_{\mathfrak{S}}(\lambda p + (1 - \lambda) q, c)$$

as required.

(iv)

$$\mathfrak{D}_{\mathfrak{S}}(p, c) - \mathfrak{D}_{\mathfrak{S}}(p, q) - \mathfrak{D}_{\mathfrak{S}}(q, c) =$$

$$\left[\sum_i p_i \mathfrak{S}(c, i) - \sum_i p_i \mathfrak{S}(p, i) \right] -$$

$$\left[\sum_i p_i \mathfrak{S}(q, i) - \sum_i p_i \mathfrak{S}(p, i) \right] -$$

$$\left[\sum_i q_i \mathfrak{S}(c, i) - \sum_i q_i \mathfrak{S}(q, i) \right] =$$

$$\sum_i (p_i - q_i) (\mathfrak{S}(c, i) - \mathfrak{S}(q, i))$$

as required.

□

Lemma 5. *Suppose \mathfrak{S} is a continuous strictly proper inaccuracy measure. Suppose \mathcal{X} is a closed convex subset of \mathcal{P} . And suppose c is not in \mathcal{X} . Then there is q in \mathcal{X} such that*

- (i) $\mathfrak{D}_{\mathfrak{S}}(q, c) < \mathfrak{D}_{\mathfrak{S}}(p, c)$ for all $p \neq q$ in \mathcal{X} .
- (ii) For all p in \mathcal{X} ,

$$\sum_{i=1}^n (p_i - q_i)(\mathfrak{S}(c, i) - \mathfrak{S}(q, i)) \geq 0$$

- (iii) For all p in \mathcal{X} ,

$$\mathfrak{D}_{\mathfrak{S}}(p, c) \geq \mathfrak{D}_{\mathfrak{S}}(p, q) + \mathfrak{D}_{\mathfrak{S}}(q, c)$$

Proof of Lemma 5. Suppose c is not in \mathcal{X} . Then, since \mathcal{X} is a closed convex set and since Lemma 4(iii) shows that $\mathfrak{D}_{\mathfrak{S}}$ is strictly convex in its first argument, there is a unique q in \mathcal{X} that minimizes $\mathfrak{D}_{\mathfrak{S}}(x, c)$ as a function of x . So, as (i) requires, $\mathfrak{D}_{\mathfrak{S}}(q, c) < \mathfrak{D}_{\mathfrak{S}}(p, c)$ for all $p \neq q$ in \mathcal{X} .

We now turn to proving (ii). We begin by observing that, since p, q are in \mathcal{P} , since \mathcal{P} is convex, and since $\mathfrak{D}_{\mathfrak{S}}(x, c)$ is minimized uniquely at $x = q$, if $0 < \varepsilon < 1$, then

$$\frac{1}{\varepsilon} [\mathfrak{D}_{\mathfrak{S}}(\varepsilon p + (1 - \varepsilon)q, c) - \mathfrak{D}_{\mathfrak{S}}(q, c)] > 0$$

Expanding that, we get

$$\frac{1}{\varepsilon} \left[\sum_i (\varepsilon p_i + (1 - \varepsilon)q_i) \mathfrak{S}(c, i) - \sum_i (\varepsilon p_i + (1 - \varepsilon)q_i) \mathfrak{S}(\varepsilon p + (1 - \varepsilon)q, i) - \sum_i q_i \mathfrak{S}(c, i) + \sum_i q_i \mathfrak{S}(q, i) \right] > 0$$

So

$$\frac{1}{\varepsilon} \left[\sum_i (q_i + \varepsilon(p_i - q_i)) \mathfrak{S}(c, i) - \sum_i (q_i + \varepsilon(p_i - q_i)) \mathfrak{S}(\varepsilon p + (1 - \varepsilon)q, i) - \sum_i q_i \mathfrak{S}(c, i) + \sum_i q_i \mathfrak{S}(q, i) \right] > 0$$

So

$$\sum_i (p_i - q_i)(\mathfrak{S}(c, i) - \mathfrak{S}(\varepsilon p + (1 - \varepsilon)q, i)) + \frac{1}{\varepsilon} \left[\sum_i q_i \mathfrak{S}(q, i) - \sum_i q_i \mathfrak{S}(\varepsilon p + (1 - \varepsilon)q, i) \right] > 0$$

Now, since \mathfrak{S} is strictly proper,

$$\frac{1}{\varepsilon} \left[\sum_i q_i \mathfrak{S}(q, i) - \sum_i q_i \mathfrak{S}(\varepsilon p + (1 - \varepsilon)q, i) \right] < 0$$

So, for all $\varepsilon > 0$,

$$\sum_i (p_i - q_i)(\mathfrak{S}(c, i) - \mathfrak{S}(\varepsilon p + (1 - \varepsilon)q, i)) > 0$$

So, since \mathfrak{S} is continuous

$$\sum_i (p_i - q_i)(\mathfrak{S}(c, i) - \mathfrak{S}(q, i)) \geq 0$$

which is what we wanted to show.

We now briefly sketch an alternative proof of (ii) that is available if $\mathfrak{D}_{\mathfrak{S}}$ is not only continuous in its first argument, but also differentiable in that argument. As before, we begin by observing that, since p, q are in \mathcal{P} , since \mathcal{P} is convex, and since $\mathfrak{D}_{\mathfrak{S}}(x, c)$ is minimized uniquely at $x = q$, then, for all $0 < \varepsilon < 1$,

$$\frac{\mathfrak{D}_{\mathfrak{S}}(\varepsilon p + (1 - \varepsilon)q, c) - \mathfrak{D}_{\mathfrak{S}}(q, c)}{\varepsilon} > 0$$

And so

$$\frac{\mathfrak{D}_{\mathfrak{S}}(q + \varepsilon(p - q), c) - \mathfrak{D}_{\mathfrak{S}}(q, c)}{\varepsilon} > 0$$

And so, if it exists,

$$\lim_{\varepsilon \rightarrow 0} \frac{\mathfrak{D}_{\mathfrak{S}}(q + \varepsilon(p - q), c) - \mathfrak{D}_{\mathfrak{S}}(q, c)}{\varepsilon} \geq 0$$

But, if it exists, the left-hand side is just the directional derivative of $\mathfrak{D}_{\mathfrak{S}}(q, c)$ with respect to q and relative to the vector $p - q$.⁸ And we know from a foundational result about directional derivatives that

$$\lim_{\varepsilon \rightarrow 0} \frac{\mathfrak{D}_{\mathfrak{S}}(q + \varepsilon(p - q), c) - \mathfrak{D}_{\mathfrak{S}}(q, c)}{\varepsilon} = \nabla \mathfrak{D}_{\mathfrak{S}}(q, c) \cdot (p - q)$$

(See (Rudin 1976, 217).) But we can also show that

$$\frac{\partial}{\partial q_k} \mathfrak{D}_{\mathfrak{S}}(q, c) = \mathfrak{S}(c, k) - \mathfrak{S}(q, k)$$

After all,

$$\begin{aligned} \frac{\partial}{\partial q_k} \mathfrak{D}_{\mathfrak{S}}(q, c) &= \frac{\partial}{\partial q_k} \left(\sum_i q_i \mathfrak{S}(c, i) - \sum_i q_i \mathfrak{S}(q, i) \right) \\ &= \mathfrak{S}(c, k) - \left(\mathfrak{S}(q, k) - \sum_i q_i \frac{\partial}{\partial q_k} \mathfrak{S}(q, i) \right) = \mathfrak{S}(c, k) - \mathfrak{S}(q, k) \end{aligned}$$

since \mathfrak{S} is strictly proper and therefore

⁸ Thanks to Jason Konek, who introduced me to the mathematics of directional derivatives in a slightly different context.

$$\sum_i q_i \frac{\partial}{\partial q_k} \mathfrak{S}(q, i) = 0$$

for all k . So

$$\begin{aligned} \nabla \mathfrak{D}_{\mathfrak{S}}(q, c) \cdot (p - q) &= \sum_i (\mathfrak{S}(c, i) - \mathfrak{S}(q, i))(p_i - q_i) \\ &= \mathfrak{D}_{\mathfrak{S}}(p, c) - \mathfrak{D}_{\mathfrak{S}}(p, q) - \mathfrak{D}_{\mathfrak{S}}(q, c) \end{aligned}$$

as required, giving us (ii).

And finally, (iii) follows immediately from (ii) and Lemma 4(iv). □

Finally, this allows us to prove Theorems 1 and 2.

Proof of Theorem 1. Suppose c is not in \mathcal{P} . Then, by Lemma 5, there is c^* in \mathcal{P} such that, for all p in \mathcal{P} ,

$$\mathfrak{D}_{\mathfrak{S}}(p, c) \geq \mathfrak{D}_{\mathfrak{S}}(p, c^*) + \mathfrak{D}_{\mathfrak{S}}(c^*, c)$$

So, in particular, since each w_i is in \mathcal{P} ,

$$\mathfrak{D}_{\mathfrak{S}}(w^i, c) \geq \mathfrak{D}_{\mathfrak{S}}(w^i, c^*) + \mathfrak{D}_{\mathfrak{S}}(c^*, c)$$

But, since c^* is in \mathcal{P} and c is not, $c^* \neq c$, and since Lemma 4(i) shows that $\mathfrak{D}_{\mathfrak{S}}$ is a divergence, $\mathfrak{D}_{\mathfrak{S}}(c^*, c) > 0$. So

$$\mathfrak{D}_{\mathfrak{S}}(w^i, c) > \mathfrak{D}_{\mathfrak{S}}(w^i, c^*)$$

So, by Lemma 4(ii), for all w_i in \mathcal{W} ,

$$\mathfrak{S}(c, i) = \mathfrak{D}_{\mathfrak{S}}(w^i, c) > \mathfrak{D}_{\mathfrak{S}}(w^i, c^*) = \mathfrak{S}(c^*, i)$$

as required. □

Proof of Theorem 2. Suppose c is not in $cl(\mathcal{A}^+)$. Then, by Lemma 5, there is c^* such that, for all p in $cl(\mathcal{A}^+)$,

$$\mathfrak{D}_{\mathfrak{S}}(p, c) \geq \mathfrak{D}_{\mathfrak{S}}(p, c^*) + \mathfrak{D}_{\mathfrak{S}}(c^*, c)$$

So, in particular, since each possible chance function ch is in $cl(\mathcal{A}^+)$,

$$\mathfrak{D}_{\mathfrak{S}}(ch, c) \geq \mathfrak{D}_{\mathfrak{S}}(ch, c^*) + \mathfrak{D}_{\mathfrak{S}}(c^*, c)$$

But, since c^* is in $cl(\mathcal{A}^+)$ and c is not, $c^* \neq c$, and since Lemma 4(i) shows that $\mathfrak{D}_{\mathfrak{S}}$ is a divergence, $\mathfrak{D}_{\mathfrak{S}}(c^*, c) > 0$. So,

$$\mathfrak{D}_{\mathfrak{S}}(ch, c) > \mathfrak{D}_{\mathfrak{S}}(ch, c^*)$$

Now,

- $\mathfrak{D}_{\mathfrak{S}}(ch, c) = \sum_i ch_i \mathfrak{S}(c, i) - \sum_i ch_i \mathfrak{S}(ch, i)$
- $\mathfrak{D}_{\mathfrak{S}}(ch, c^*) = \sum_i ch_i \mathfrak{S}(c^*, i) - \sum_i ch_i \mathfrak{S}(ch, i)$,

so

$$\sum_i ch_i \mathfrak{S}(c, i) > \sum_i ch_i \mathfrak{S}(c^*, i)$$

as required. □

6.2. Proof of Theorem 3

To prove Theorem 3, we need a divergence not between one credence function and another, but between a sequence of $m + 1$ credence functions and another sequence of $m + 1$ credence functions. We create that in the natural way. That is, given p^0, p^1, \dots, p^m in \mathcal{P} and c^0, c^1, \dots, c^m in \mathcal{C} , the divergence from the former sequence to the latter is just the sum of the divergences from p^0 to c^0 , p^1 to c^1 , and so on. Thus:

Corollary 6. *Suppose \mathfrak{S} is a strictly proper inaccuracy measure. Then define $\mathfrak{D}_{\mathfrak{S}} : \mathcal{P}^{m+1} \times \mathcal{C}^{m+1} \rightarrow [0, \infty]$ as follows:*

$$\mathfrak{D}_{\mathfrak{S}}((p^0, p^1, \dots, p^m), (c^0, c^1, \dots, c^m)) := \sum_{k=0}^m \sum_{i=1}^n \left(p_i^k \mathfrak{S}(c^k, i) - \sum_i p_i^k \mathfrak{S}(p^k, i) \right)$$

Then:

- (i) $\mathfrak{D}_{\mathfrak{S}}$ is a divergence.
- (ii) $\mathfrak{D}_{\mathfrak{S}}((w^i, c^1, \dots, c^{k-1}, w^i, c^{k+1}, \dots, c^m), (c^0, c^1, \dots, c^m)) = \mathfrak{S}(c^0, i) + \mathfrak{S}(c^k, i)$, for all $1 \leq k \leq m$ and $1 \leq i \leq n$.
- (iii) $\mathfrak{D}_{\mathfrak{S}}$ is strictly convex in its first argument.

Proof of Corollary 6. These follow immediately from lemma 4. □

Corollary 7. *Suppose \mathfrak{S} is a continuous strictly proper inaccuracy measure. Suppose \mathcal{X} is a closed convex subset of \mathcal{P}^{m+1} . And suppose (c^0, c^1, \dots, c^m) is not in \mathcal{X} . Then there is (q^0, q^1, \dots, q^m) in \mathcal{X} such that*

- (i) $\sum_{k=0}^m \mathfrak{D}_{\mathfrak{S}}(q^k, c^k) < \sum_{k=0}^m \mathfrak{D}_{\mathfrak{S}}(p^k, c^k)$ for all $(p^0, p^1, \dots, p^m) \neq (q^0, q^1, \dots, q^m)$ in \mathcal{X} .
- (ii) For all (p^0, p^1, \dots, p^m) in \mathcal{X} ,

$$\sum_{k=0}^m \left(\sum_{i=1}^n (p_i^k - q_i^k) (\mathfrak{S}(c^k, i) - \mathfrak{S}(q^k, i)) \right) \geq 0$$

- (iii) For all (p^0, p^1, \dots, p^m) in \mathcal{X} ,

$$\sum_{k=0}^m \mathfrak{D}_{\mathfrak{S}}(p^k, c^k) \geq \sum_{k=0}^m \mathfrak{D}_{\mathfrak{S}}(p^k, q^k) + \sum_{k=0}^m \mathfrak{D}_{\mathfrak{S}}(q^k, c^k)$$

Proof of Corollary 7. The proof strategy is exactly as for Lemma 5. □

To prove theorem 3, we now need just one more result:

Lemma 8. Given c^0, c^1, \dots, c^m in \mathcal{P} , let

$$\mathcal{X} = \{(w^i, c^1, \dots, c^{k-1}, w^i, c^{k+1}, \dots, c^m) : w_i \in \mathcal{W} \quad 1 \leq k \leq m\}$$

Then,

(i) $\mathcal{X}^+ \subseteq \mathcal{P}^{m+1}$.

(ii) If c^0 is not in the convex hull of c^1, \dots, c^m , then (c^0, c^1, \dots, c^m) is not in \mathcal{X}^+ .

Proof of Lemma 8. \mathcal{P} is closed under taking mixtures, which gives us (i). We prove (ii) by proving the contrapositive. Suppose (c^0, c^1, \dots, c^m) is in \mathcal{X}^+ . Then there are $0 \leq \lambda_{i,k} \leq 1$ such that $\sum_{i=1}^n \sum_{k=1}^m \lambda_{i,k} = 1$ and

$$(c^0, c^1, \dots, c^m) = \sum_{i=1}^n \sum_{k=1}^m \lambda_{i,k} (w^i, c^1, \dots, c^{k-1}, w^i, c^{k+1}, \dots, c^m)$$

Thus,

$$c^0 = \sum_{i=1}^n \sum_{k=1}^m \lambda_{i,k} w^i$$

and

$$c^k = \sum_{i=1}^n \lambda_{i,k} w^i + \sum_{i=1}^n \sum_{l \neq k} \lambda_{i,l} c^k$$

So

$$\left(\sum_{i=1}^n \lambda_{i,k} \right) c^k = \sum_{i=1}^n \lambda_{i,k} w^i$$

So let $\lambda_k = \sum_{i=1}^n \lambda_{i,k}$. Then, for $1 \leq k \leq m$,

$$\lambda_k c^k = \sum_{i=1}^n \lambda_{i,k} w^i$$

And thus

$$\sum_{k=1}^m \lambda_k c^k = \sum_{k=1}^m \sum_{i=1}^n \lambda_{i,k} w^i = c^0$$

as required.

Now we can turn to the proof of Theorem 3. □

Proof of Theorem 3. If c^0 is not in the convex hull of c^1, \dots, c^m , then (c^0, c^1, \dots, c^m) is not in \mathcal{X}^+ . Thus, by Lemma 7, there is (q^0, q^1, \dots, q^m) such that, for all (p^0, p^1, \dots, p^m) in \mathcal{X}^+ ,

$$\mathfrak{D}_{\mathfrak{S}}((p^0, p^1, \dots, p^m), (q^0, q^1, \dots, q^m)) < \mathfrak{D}_{\mathfrak{S}}((p^0, p^1, \dots, p^m), (c^0, c^1, \dots, c^m))$$

In particular, for w_i in \mathcal{W} and $1 \leq k \leq m$,

$$\begin{aligned} & \mathfrak{D}_{\mathfrak{S}}((w^i, c^1, \dots, c^{k-1}, w^i, c^{k+1}, \dots, c^m), (q^0, q^1, \dots, q^m)) \\ & < \mathfrak{D}_{\mathfrak{S}}((w^i, c^1, \dots, c^{k-1}, w^i, c^{k+1}, \dots, c^m), (c^0, c^1, \dots, c^m)) \end{aligned}$$

But

$$\begin{aligned} \mathfrak{S}(q^0, i) + \mathfrak{S}(q^k, i) &= \mathfrak{D}_{\mathfrak{S}}(w^i, q^0) + \mathfrak{D}_{\mathfrak{S}}(w^i, q^k) \\ &\leq \mathfrak{D}_{\mathfrak{S}}((w^i, c^1, \dots, c^{k-1}, w^i, c^{k+1}, \dots, c^m), (q^0, q^1, \dots, q^m)) \\ &< \mathfrak{D}_{\mathfrak{S}}((w^i, c^1, \dots, c^{k-1}, w^i, c^{k+1}, \dots, c^m), (c^0, c^1, \dots, c^m)) \\ &= \mathfrak{D}_{\mathfrak{S}}(w^i, c^0) + \mathfrak{D}_{\mathfrak{S}}(w^i, c^k) \\ &= \mathfrak{S}(c^0, i) + \mathfrak{S}(c^k, i) \end{aligned}$$

as required. □