Graded Incoherence for Accuracy-Firsters

Glauber De Bona and Julia Staffel*†

This article investigates the relationship between two evaluative claims about agents' degrees of belief: (i) that it is better to have more rather than less accurate degrees of belief and (ii) that it is better to have less rather than more probabilistically incoherent degrees of belief. We show that, for suitable combinations of inaccuracy measures and incoherence measures, both claims are compatible, although not equivalent; moreover, certain ways of becoming less incoherent always guarantee improvements in accuracy. Incompatibilities between particular incoherence and inaccuracy measures can be exploited to argue against particular ways of measuring either inaccuracy or incoherence.

1. Introduction. Proponents of accuracy-first epistemology defend the view that the most fundamental epistemic end is to have degrees of belief that are as accurate (i.e., as close to the truth) as possible.[1] Other epistemic norms or goods are only valuable to the extent that they help promote the fundamental end of having accurate credences. Accuracy-first epistemology is attractive because many commonly endorsed epistemic norms are in fact derivable from the end of accuracy.[2] One of the central results in accuracy-first epistemology is the accuracy argument for probabilism (i.e., the claim

Received June 2015; revised March 2016.

*To contact the authors, please write to: Glauber De Bona, University College London; e-mail: glauberbona@gmail.com. Julia Staffel, Washington University in St. Louis; e-mail: staffel.julia@gmail.com.

†Glauber De Bona is financially supported by CNPq (National Council for Scientific and Technological Development). We would like to thank the reviewers, Leszek Wronski, Liam Kofi Bright, and Carlos Stein Naves de Brito for their valuable insights. Both authors contributed equally.

1. Goldman defends an analogous view about outright belief called 'veritism' (2002, 52).

2. Philosophers who have argued for or against epistemic norms on accuracy-based grounds include, among many others, Joyce (1998, 2009), Greaves and Wallace (2006), Moss (2011), Williams (2012), Caie (2013), Easwaran (2013), Staffel (2015a), and Pettigrew (2016).

Philosophy of Science, 84 (April 2017) pp. 189–213. 0031-8248/2017/8402-0001$10.00
Copyright 2017 by the Philosophy of Science Association. All rights reserved.

https://doi.org/10.1086/690715 Published online by Cambridge University Press

that rational agents should have probabilistically coherent rather than incoherent degrees of belief).

It has recently been argued that we should understand incoherence as a graded notion. Credence functions can be more or less incoherent, depending on how well they approximate a probability function (Zynda 1996; Schervish, Seidenfeld, and Kadane 2000, 2002; Schervish, Kadane, and Seidenfeld 2003; Wang et al. 2011; De Bona and Finger 2015; Staffel 2015b). This invites the natural thought that incoherent credence functions should not all be treated the same from a normative point of view. It is plausible to assume that it is better to be less rather than more incoherent because credence functions seem epistemically better, the more they approximate the ideal of probabilistic coherence. However, the accuracy argument for probabilism is not concerned with differentiating between incoherent credence functions; its aim is to show that all incoherent credence functions, unlike coherent ones, are epistemically defective.

Since accuracy-firsters seek accuracy-based justifications of independently plausible epistemic norms, we ask whether they can support the claim that it is better to be less rather than more incoherent. We settle this question by investigating how the claim that less incoherence is better relates to the accuracy-firsters' claim that less inaccuracy is better. We show that, for suitable combinations of inaccuracy measures and incoherence measures, both principles are compatible, although not equivalent; moreover, certain ways of becoming less incoherent always guarantee improvements in accuracy. We also show how incompatibilities between particular incoherence and inaccuracy measures can be exploited to argue against particular ways of measuring either inaccuracy or incoherence.

In section 2, we present the familiar accuracy-based argument for probabilism. In section 3, we introduce the idea that incoherence is a graded notion and explain how degrees of incoherence can be measured. In section 4, we explore to what extent the claim that it is better to be less rather than more incoherent is compatible with the view that accuracy is the ultimate source of epistemic value. In section 5, we summarize our results and point to avenues for further research. The proofs can be found in the appendix.

2. Justifying Epistemic Norms in Accuracy-First Epistemology. Accuracy-firsters rely on a basic argumentative strategy to establish normative epistemological claims. They begin by defining what epistemic value, or epistemic utility, consists in. Then they plug this definition into a decision theoretic principle and prove that a particular class of epistemic states is favored because these states are not worse than any alternative state at attaining epistemic value regarding that principle.

Accuracy-firsters use this strategy to argue for probabilism, which is the claim that a rational agent should have a probabilistically coherent credence function. A credence function is a function $c : F \to [0, 1]$ that assigns a real

https://doi.org/10.1086/690715 Published online by Cambridge University Press

number from the interval [0, 1] to each proposition in a finite set of propositions $F \subseteq \mathcal{L}$.[3] The language $\mathcal{L}$ is the set of propositions generated by taking a finite set of atomic statements $\{A_i\}$ and combining it with the standard Boolean connectives $\neg, \vee, \wedge$. A probabilistic credence function can then be defined as follows: a function $P : \mathcal{L} \rightarrow \mathbb{R}$ is a probability function on $\mathcal{L}$ if it obeys the axioms:

i) For any proposition $\varphi \in \mathcal{L}$, if φ is a tautology, then $P(\varphi) = 1$,
ii) For any proposition $\varphi \in \mathcal{L}$, $P(\varphi) \geq 0$.
iii) For any propositions $\varphi, \psi \in \mathcal{L}$, if φ and ψ are mutually exclusive, then $P(\varphi \vee \psi) = P(\varphi) + P(\psi)$.

For a credence function $c : F \rightarrow [0, 1]$ to be probabilistically coherent, there must be a probability function $P : \mathcal{L} \rightarrow [0, 1]$ such that $P(\varphi) = c(\varphi)$ for all $\varphi \in F$.

The argument for probabilism then goes like this: first, we define an optimal or maximally accurate credence function at a given world. Accuracy-firsters propose that the best credence to have in some proposition φ at world w is 1 if φ is true at w, and 0 if φ is false at w.

Next, we need a measure of inaccuracy: if an agent's credences are not maximally accurate, we need to measure their distance from the truth in a given world. For the argument to work, we can choose any inaccuracy measure that is based on a continuous, strictly proper scoring rule.[4]

Next, we need a principle of decision theory to establish the claim that only probabilistic credences are rational. Suppose that we have a set O of options $\{o_1, \ldots, o_n\}$ and a set of worlds W. Let U be a utility function, such that U takes an option o_j from O and a world w_i from W and returns a real number that gives the utility of o_j at w_i. An option o_j strongly dominates an option o_k just in case the utility of o_j is greater than the utility of o_k in every possible world. An option o_j weakly dominates an option o_k just in case the utility of o_j is at least as great as the utility of o_k in every possible world and strictly greater in at least one world.

This helps us state a sufficient condition for an option to be irrational for an agent, relative to U:[5]

3. We take it to be a definitional truth that credences must lie in the interval [0, 1]. We could in principle choose any other interval $[x, y]$ and designate x as the value given to the lowest possible credence and y as the value given to the highest possible credence.

4. A popular inaccuracy measure in the literature is the Brier score. Let $c : F \rightarrow [0, 1]$ be a credence function, and let I_w be the indicator function that returns 1 if some proposition φ is true at w, and 0 if it is false at w. Then, the Brier score of c at world w is $B(c, w) = \Sigma_{\varphi \in F}(c(\varphi) - I_w(\varphi))^2$. Pettigrew (2016) argues for adopting the Brier score as the best scoring rule.

5. As an anonymous referee points out, one could remove the clause that says "o_k is not weakly dominated by any other option relative to U," which would make any dominated

https://doi.org/10.1086/690715 Published online by Cambridge University Press

Undominated Dominance: If some option o_j is strongly dominated by another option o_k relative to U, and o_k is not weakly dominated by any other option relative to U, then option o_j is irrational for an agent whose utility function is U.

These assumptions entail that there is a significant difference between probabilistically coherent and probabilistically incoherent credence functions. If a credence function is incoherent, then there is an alternative coherent credence function whose inaccuracy is lower at every possible world and that is not itself weakly dominated by another credence function. By contrast, if a credence function is coherent, then there is no alternative credence function that even weakly accuracy-dominates it. Hence, the dominance principle rules out incoherent credence functions as irrational, but it does not rule out coherent credence functions.[6] The accuracy-firster thus succeeds in showing that an agent whose ultimate goal is to have accurate credences should have coherent rather than incoherent credences. On this view, coherence is not an independent epistemic end; rather, it is valuable because it promotes accuracy.

3. Coherence as a Graded Notion. The accuracy argument for probabilism shows that coherent credence functions have a desirable epistemic property that incoherent credence functions lack. Given its focus, it does not pay attention to differences between incoherent credence functions. Yet, intuitively, not all incoherent credence functions defined on the same set of propositions are equally incoherent. For example, suppose Jane has a credence of 0.2 in some tautology $\top$, whereas Clara's credence is 0.99. Intuitively, Jane is more incoherent than Clara, given that the credence in $\top$ required to be coherent is 1. Or, consider a case in which Clara and Jane each have credences in a partition of propositions $\{\varphi, \psi, \theta\}$. Clara assigns each proposition in the partition a credence of 0.35. Jane assigns each proposition a credence of 0.25. Each of them has incoherent credences because their credences in φ, ψ, and θ do not sum to 1. Yet, it is again natural to think that Clara's credences are less incoherent than Jane's because they seem closer to the ideal of being coherent.

Thinking about incoherence as a graded notion raises at least two questions, which we will take up in the following two subsections: first, can we

options irrational. The reason for including this clause is that there are cases involving infinitely many options, each of which is dominated. If we do not want every possible option to be irrational in these cases, we need to keep this clause. For a good discussion of this, see Pettigrew (2016).

6. Various versions of this argument can be found in the literature; see De Finetti (1974), Joyce (1998, 2009), Predd et al. (2009), and Pettigrew (2016).

https://doi.org/10.1086/690715 Published online by Cambridge University Press

find a general way of measuring degrees of incoherence that makes those pretheoretic intuitions precise? Second, is there any normative significance to being more or less incoherent?

3.1. How Can We Measure Degrees of Incoherence? With regard to the first question, we can find numerous proposals for measures of incoherence in the literature. They fall into three categories: qualitative incoherence measures, Dutch book measures, and distance measures or divergences. A measure of the first category is proposed by Zynda (1996). His measure ranks incoherent credence functions by comparing their largest coherent subsets. A class of measures in the second category has been proposed by Schervish et al. (2000, 2002, 2003). Their idea is to exploit the fact that we can make Dutch books against incoherent agents: we can set up a collection of bets, each of which is justified by one of the agent's credences but that together guarantee that the agent loses money. If we normalize these Dutch book losses, then we can measure incoherence by checking the rate at which an incoherent agent is bound to lose money. Different normalizations generate different measures of incoherence. The third class of measures, which has been explored by De Bona and Finger (2015) and used in research by Wang et al. (2011), measures the incoherence of a credence function by determining its divergence from some closest coherent credence function. A credence function is represented as a vector, and thus the task is to determine its divergence from some closest vector that represents a coherent credence function. Different incoherence measures can be generated by selecting different divergences.

In statistics, a divergence is a nonnegative function d that measures the "distance" between two probability distributions, say p and q, where $d(\mathrm{p}, q) = 0$ if and only if (iff) $p = q$. More generally, to handle credence functions on arbitrary sets of propositions, we define a divergence as any function $d : [0, 1]^n \times [0, 1]^n \rightarrow [0, \infty)$, over vectors $X, Y \in [0, 1]^n$, where $d(X, Y) = 0$ iff $X = Y$. A subset of divergences are metrics, which are the kind of measures we usually associate with distance measures and which additionally have the properties of being symmetrical and satisfying the triangle inequality. The p-norm distance between two vectors $X = \langle x_1, x_2, \ldots, x_n \rangle$ and $Y = \langle y_1, y_2, \ldots, y_n \rangle$ is defined as the p-norm of their difference, with $p \in [1, \infty]$:

$$\| X - Y \|_p = \left(\sum_{i=1}^{n} |x_i - y_i|^p \right)^{1/p}.$$

For example, we can set $p = 1$, which gives us the absolute distance measure, or $p = 2$, which gives us Euclidean distance, and so on.

We can apply divergences to credence functions on the same set of propositions through their corresponding vectors. Given a credence function

https://doi.org/10.1086/690715 Published online by Cambridge University Press

$c : \{\varphi_1, \ldots, \varphi_m\} \to [0, 1]$ such that $c(\varphi_i) = q_i$ for $1 \leq i \leq m$, we say that $q = \langle q_1, \ldots, q_m \rangle$ is its corresponding vector. Abusing the notation, we also use c to represent the corresponding vector where the context avoids ambiguity. Now d can be used to formulate an incoherence measure $\mathcal{I}_d$, defined for each credence function c as the minimum divergence to a coherent credence function c':

$$\mathcal{I}_d(c) = \min\{d(c, c') | c' \text{ is coherent}\}.$$

We say the incoherence measure $\mathcal{I}_d$ is generated by the divergence d. Every such incoherence measure is nonnegative and has the additional desirable property, usually called *consistency* (De Bona and Finger 2015), that, for any credence function c, $\mathcal{I}_d(c) = 0$ iff c is coherent.

Which of these incoherence measures should we focus on? Staffel (2015b) has argued that we should focus on measuring what she calls the 'overall incoherence' of a credence function. Such a measure tracks changes in incoherence that appear anywhere in the agent's credences. A measure that is suitable for this task must avoid relying on an unrepresentative sample of the agent's credences because it might thereby miss differences in incoherence. For example, Schervish, Seidenfeld, and Kadane's neutral/sum measure is ruled out by this criterion because it determines the incoherence of a credence function based only on the subset of the credence function that leads to the worst loss ratio in a Dutch book. It is insensitive to whether any incoherence occurs in the agent's credences outside of this worst-case subset. By contrast, the p-norm distance measures of incoherence that can be generated with a finite value for p avoid this problem, since they do not single out some subset of the agent's credences as the determinant of incoherence. For this reason, Staffel argues that they are especially well suited for measuring overall incoherence.[7] These p-norm distance measures determine incoherence by representing the entire credence function as a vector and measuring the distance to some closest vector that represents a coherent credence function in a way that does not disregard any parts of the agent's credence function.

On the basis of Staffel's arguments, we focus on the third class of measures. While Staffel does not discuss divergences that are not metrics, we will take into consideration all incoherence measures that are based on continuous divergences. Focusing on continuous divergences ensures that the

7. The motivation to rule out the measure where $p = \infty$, called the Chebyshev distance, is that the corresponding incoherence measure is equivalent to the neutral/sum Dutch book measure proposed by Schervish, Seidenfeld, and Kadane, which Staffel argues to be unsuitable for measuring overall incoherence. This equivalence and further results about the relationship between Dutch book measures and p-norm distance measures of incoherence are proved in De Bona and Finger (2015).

https://doi.org/10.1086/690715 Published online by Cambridge University Press

distance to some closest coherent credence function is well defined for any finite credence function.[8] This aligns our selection of incoherence measures with a further constraint proposed by Staffel, according to which a suitable measure of overall incoherence should be applicable to arbitrary finite credence functions and make them comparable. Moreover, we favor measuring incoherence based on divergences rather than metrics because it helps us bring out crucial connections between incoherence and inaccuracy. Each additive inaccuracy measure based on a continuous, strictly proper scoring rule is naturally associated with a particular divergence. As we discuss in section 4, we can prove some attractive results about the relationship between accuracy and coherence if we pair inaccuracy measures with incoherence measures that are based on their associated divergence.

To illustrate how distance/divergence-based incoherence measures work, we will apply the absolute distance measure to one of our examples. Considering the partition $\{\varphi, \psi, \theta\}$, Clara's credences are $\mathrm{Cr}_C(\varphi) = \mathrm{Cr}_C(\psi) = \mathrm{Cr}_C(\theta) = 0.35$. The minimum distance to a coherent credence function is 0.05 for Clara, where $P(\varphi) = 0.35, P(\psi) = 0.35, P(\theta) = 0.3$ is an example of a coherent credence function that minimizes the distance to Clara's incoherent credences. Jane's credences are $\mathrm{Cr}_J(\varphi) = \mathrm{Cr}_J(\psi) = \mathrm{Cr}_J(\theta) = 0.25$. The minimum distance to a coherent credence function is 0.25 for Jane, with $P(\varphi) = 0.35$, $P(\psi) = 0.35$, $P(\theta) = 0.3$ also being such a minimally distant credence function for Jane. In sum, we can answer our first question affirmatively: it is possible to give a formally precise account of what it means for one credence function to be more incoherent than another.

3.2. Is There Any Normative Significance to Being More or Less Incoherent? The second question concerns the normative significance of degrees of incoherence. Is there a difference between incoherent credence functions, in the sense that it is worse to have more rather than less incoherent credences and, equivalently, better to have less rather than more incoherent credences? We have already seen that even if we focus just on incoherence measures based on continuous divergences, there are many nonequivalent ways of measuring incoherence. Hence, a defender of the claim that less incoherence is better must specify how degrees of incoherence should be determined. Here is a first schematic formulation of this claim, to be refined later. Let $\mathcal{I}_d$ be a measure of incoherence, which is generated by a divergence d. Then, we can generate different instances of the claim that less incoherence is better by filling in different values for $\mathcal{I}_d$.

8. As the set of vectors corresponding to coherent credence functions (a convex hull) is compact, any continuous function on it has bounded, closed range, yielding a well-defined minimum.

https://doi.org/10.1086/690715 Published online by Cambridge University Press

Graded Incoherence Principle: For any credence functions $c : F \to [0, 1]$ and $c' : F \to [0, 1]$, if c is less $\mathcal{I}_d$-incoherent than c', then c is epistemically better than c'.

In formal epistemology, probabilistic coherence is usually treated as a rational ideal. Yet, this ideal is out of reach for ordinary human agents because of their computational limitations (see, e.g., Hacking 1967; Harman 1986). But how can an unattainable ideal still have normative significance for ordinary agents? Zynda (1996) and Staffel (2015b) argue that if an end is unattainable for an agent, it can still be relevant for evaluating or guiding the agent if there is a meaningful way of characterizing approximations to the end. By contrast, if an end is unattainable, and cannot be approximated, then it is hard to see how it could serve an interesting guiding or evaluative function. The inability of limited agents to perfectly conform to normative ideals thus justifies an interest in normative claims that concern the comparative value of nonideal ways of being. We want to know whether, for example, a change in a nonideal agent's attitudes constitutes an improvement or a deterioration, even if neither the initial nor the subsequent state conforms to normative ideals.

This is exactly how accuracy-firsters think about the fundamental end of accuracy. The ideal state is to have maximally accurate credences, or in other words, to be omniscient. This state is not attainable for ordinary agents. Still, accuracy can serve as an end for ordinary agents because they can improve the inaccuracy of their credences, and we can use inaccuracy measures to track changes in their credences' inaccuracy. By contrast, it is hard to see how the end of having maximally accurate credences could play this role if either there was no way of measuring degrees of inaccuracy or there was no sense in which having more inaccurate credences was worse than having less inaccurate credences.

We can think of coherence on the same model. There are both theoretical and empirical reasons to think that there are normative differences between incoherent credence functions. Agents whose credences are incoherent have a disunified view of the world. This does not seem epistemically good, and it can also lead the agent to act in self-undermining ways. Dutch books are just a somewhat artificial illustration of what happens when an agent's credences are incoherent. If we hold betting stakes fixed, becoming more incoherent leads to greater losses. Of course in real life, there need not be a perfect correspondence between how incoherent an agent's credences are and how badly the incoherence affects the unity of the agent's plans and actions. Yet, it is still plausible that being less incoherent is generally better. An interesting empirical result in this context comes from a study by Wang et al. (2011), who conducted a survey about the outcome of the 2008 US presidential elections. Over 15,000 subjects were asked to give probabilistic

https://doi.org/10.1086/690715 Published online by Cambridge University Press

forecasts of the election results in various states. The accuracy of people's forecasts was measured by the Brier score, and the incoherence of their forecasts was measured by the squared Euclidean distance to the closest coherent forecast (although the authors point out that their results still hold if incoherence is measured on the basis of absolute distance). People tended to give incoherent forecasts, but those who gave less incoherent forecasts were more accurate predictors of the actual outcome of the election than people who gave more incoherent forecasts. This shows that being less incoherent is at least correlated with other desirable epistemic features, such as making accurate predictions.[9]

How exactly we want to spell out the Graded Incoherence Principle depends on the role the coherence norm plays. First, one might take a coherence-first position, which takes having coherent credences to be the fundamental epistemic aim. Coherence-firsters only endorse other epistemic norms insofar as they promote the end of being coherent. Of course, a coherence-firster would have to decide on an appropriate measure of degrees of incoherence (just like an accuracy-firster must decide on an inaccuracy measure), but once a measure has been determined, the coherence-firster is committed to claiming that being less incoherent according to this measure is epistemically better than being more incoherent. Moreover, whenever there might be reasons for being more incoherent, they are overridden by coherence-promoting norms according to this view because coherence is the most fundamental epistemic value.

Second, one might be a pluralist about epistemic norms. For example, one might consider the coherence norm to be an epistemic norm but also recognize other norms as being equally fundamental, such as the norm to proportion one's credences to the strength of one's evidence (see, e.g., Christensen [2007] for such a view). A pluralist can endorse a version of the Graded Incoherence Principle, although conflicts with other epistemic norms might occur. On the pluralist view, recognizing the possibility of norm conflicts is compatible with considering an instance of the Graded Incoherence Principle to be a genuine epistemic norm.

Finally, one might consider coherence to be a derivative norm, which is the accuracy-firsters' position. Accuracy-firsters only endorse norms that promote accuracy. Moreover, they want to derive as many commonly endorsed epistemic norms as possible from the end of accuracy because doing so supports their view that accuracy really is the fundamental epistemic end. Thus,

9. As an anonymous referee points out, one possible explanation of this phenomenon is that less incoherent people tend to be smarter, which is why they are also likely to be more accurate. Of course, this explanation already suggests that coherence is desirable and that smarter agents are better at approximating this desirable end.

https://doi.org/10.1086/690715 Published online by Cambridge University Press

accuracy-firsters will only endorse a version of the Graded Incoherence Principle if it can be underwritten on accuracy-based grounds.

Being able to defend a version of this principle would be attractive for the accuracy-firster for two reasons: first, this would give the accuracy-firster another piece of evidence for her claim that accuracy really is the fundamental epistemic end. Second, it might prevent the accuracy-firster from having to adopt an unattractive position regarding coherence. Suppose her view that accuracy is the only epistemic value is incompatible with the claim that it is better to be less incoherent than to be more incoherent, regardless of what incoherence measure is chosen. Then she would have to claim that ideal rationality requires perfect coherence and claim at the same time that there is no general recommendation that agents should try to approximate this ideal because nonideal agents are sometimes better off moving away from it. While this is a consistent position the accuracy-firster could hold, it is certainly not as attractive as the alternative. Hence, it would be desirable for the accuracy-firster if we could derive a version of the Graded Incoherence Principle from the fundamental goal of accuracy or at least show their compatibility.

4. The Graded Incoherence Principle for Accuracy-Firsters. Accuracy-firsters think that it is better to have more rather than less accurate credences. In order to find out how their view is related to the Graded Incoherence Principle, we must make their position formally precise. We will assume, as before, that the most accurate credence function at a given world assigns credence 1 to all the truths and credence 0 to all the falsehoods. Fix a set $F = \{\varphi_1, \dots, \varphi_m\}$ of formulas, and let $V \subset \{0, 1\}^m$ denote the set of vectors corresponding to the possible worlds (i.e., to the most accurate credence function at these worlds). An inaccuracy measure $\mathfrak{I} : [0, 1]^m \times V \to \mathbb{R}$ is a function that takes the vectors corresponding to a credence function and to a possible world and returns a real value. Such a function is said to be additive if there is a scoring rule $s : [0, 1] \times \{0, 1\} \to \mathbb{R}$ such that, for all $q \in [0, 1]^m$ and $v \in V$, $\mathfrak{I}(q, v) = \Sigma_{i=1}^m s(q_i, v_i)$. In this case, we say $\mathfrak{I}$ is generated by the scoring rule s. We say a scoring rule s is strictly proper when, for any $\eta \in [0, 1]$, only $x = \eta$ minimizes $\eta s(x, 1) + (1 - \eta)s(x, 0)$ and focus on continuous inaccuracy measures generated by strictly proper scoring rules.

We will understand accuracy-dominance relations as defined in section 2, taking credence functions on a fixed $F \subseteq \mathcal{L}$ as options and accuracy as utility. Here are then two principle schemata from which we can generate accuracy-dominance principles by specifying an inaccuracy measure $\mathfrak{I}$. The accuracy-firster will endorse the versions of the principles that employ her preferred inaccuracy measure since she thinks that it is better to have more $\mathfrak{I}$-accurate credences.

https://doi.org/10.1086/690715 Published online by Cambridge University Press

Weak Accuracy-Dominance Principle: For any credence functions $c : F \rightarrow [0, 1]$ and $c' : F \rightarrow [0, 1]$, if c weakly $\mathfrak{I}$-accuracy-dominates c', then c is epistemically better than c'.

Strong Accuracy-Dominance Principle: For any credence functions $c : F \rightarrow [0, 1]$ and $c' : F \rightarrow [0, 1]$, if c strongly $\mathfrak{I}$-accuracy-dominates c', then c is epistemically better than c'.

We can now investigate how these principles relate to the Graded Incoherence Principle, which we also stated in schematic form in the previous section:

Graded Incoherence Principle: For any credence functions $c : F \rightarrow [0, 1]$ and $c' : F \rightarrow [0, 1]$, if c is less $\mathcal{I}_d$-incoherent than c', then c is epistemically better than c'.

In principle, there are four different (but not mutually exclusive) relations that can hold between two epistemic principles A and B.

1. A is equivalent to B: for any credence functions c and c', c is epistemically better than c' according to A iff it is according to B;
2. A implies B (B implies A): for any credence functions c and c', if c is epistemically better than c' according to B (A), then it is according to A (B);
3. A is compatible with B: there is no pair of credence functions c and c' such that c is epistemically better than c' according to A and c' is epistemically better than c according to B;
4. A is incompatible with B: there is some pair of credence functions c and c' such that c is epistemically better than c' according to A and c' is epistemically better than c according to B.

We are interested in investigating how instances of the Accuracy-Dominance Principles relate to instances of the Graded Incoherence Principle. We will do so by stating a series of schematic claims and checking whether we can find any true instances of them by substituting pairs of incoherence measures and inaccuracy measures for $\mathcal{I}_d$ and $\mathfrak{I}$. To keep track of which claims have true instances and which ones do not, we will mark the claims that do not have true instances with an asterisk. We will begin by investigating whether we can find any equivalence relations. These relations are expressed in claims Ia* and Ib*:

(Ia*) *Reducing incoherence promotes accuracy, and promoting accuracy reduces incoherence* (*weak dominance version*): For all credence func-

https://doi.org/10.1086/690715 Published online by Cambridge University Press

tions c, c' defined on the same set of propositions, c is less $\mathcal{I}_d$-incoherent than c' iff c weakly $\mathfrak{I}$-accuracy-dominates c'.

(Ib*) *Reducing incoherence promotes accuracy, and promoting accuracy reduces incoherence* (*strong dominance version*): For all credence functions c, c' defined on the same set of propositions, c is less $\mathcal{I}_d$-incoherent than c' iff c strongly $\mathfrak{I}$-accuracy-dominates c'.

Claim Ia* states an equivalence schema between the Graded Incoherence Principle and the Weak Accuracy-Dominance Principle, and claim Ib*, between the former and the Strong Accuracy-Dominance Principle. We can rule out both Ia* and Ib* for any combinations of incoherence measures $\mathcal{I}_d$ and inaccuracy measures $\mathfrak{I}$. According to every incoherence measure $\mathcal{I}_d$, every coherent credence function is less $\mathcal{I}_d$-incoherent than every incoherent credence function. However, there is no inaccuracy measure—based on a continuous, strictly proper scoring rule—for which it holds that every coherent credence function $\mathfrak{I}$-accuracy-dominates every incoherent credence function. Hence, every instance of the schemata Ia* and Ib* has counterexamples (see proposition 2 in the appendix).

Here is an example that illustrates this finding: take the incoherent credence function $c(\varphi) = 0.6$, $c(\neg\varphi) = 0.6$, for some atomic proposition φ. The coherent credence function $c^*(\varphi) = 0.9$, $c^*(\neg\varphi) = 0.1$ is less incoherent than c on any incoherence measure $\mathcal{I}_d$. Yet, if we measure inaccuracy with the Brier score, c^* does not even weakly dominate c. We can find suitable counterexamples of this kind regardless of which inaccuracy measure we choose.

These counterexamples show that the left-to-right directions of the biconditionals have no true instances; hence, we can also rule out the following schemata IIa* and IIb*.

(IIa*) *Reducing incoherence promotes accuracy* (*weak dominance version*): For all credence functions c, c' defined on the same set of propositions, if c is less $\mathcal{I}_d$-incoherent than c', then c weakly $\mathfrak{I}$-accuracy-dominates c'.

(IIb*) *Reducing incoherence promotes accuracy* (*strong dominance version*): For all credence functions c, c' defined on the same set of propositions, if c is less $\mathcal{I}_d$-incoherent than c', then c strongly $\mathfrak{I}$-accuracy-dominates c'.

Hence, not every credence function that is less incoherent is also epistemically better on accuracy-based grounds. Yet, it is still possible that the accuracy-firster can endorse a suitably qualified version of the idea that it is better to be less rather than more incoherent. This just shows that she can-

https://doi.org/10.1086/690715 Published online by Cambridge University Press

not underwrite the idea that, if your credences are incoherent, any less incoherent credence function is epistemically better. But this claim seems very strong anyway. While a coherence-firster might endorse it, a pluralist about epistemic norms would probably deny it and claim instead that other norms, for example, the norm about proportioning your credences to your evidence, constrain the ways in which an incoherent agent should reduce her incoherence. Hence, an accuracy-firster would not be alone in thinking that reducing incoherence is only good when it is done in the right way. Maybe some weaker form of the Graded Incoherence Principle is implied by an Accuracy-Dominance Principle, for suitable choices of d—consequently, of $\mathcal{I}_d$—and $\mathfrak{I}$. We will investigate this later, after analyzing the other direction of the implication.

We now examine the remaining two possibilities of implication: whether the Graded Incoherence Principle implies some Accuracy-Dominance Principle, for suitable choices of $\mathcal{I}_d$ and $\mathfrak{I}$. These entailments, if correct, would not justify the accuracy-firsters' support for the Graded Incoherence Principle, but they would imply that supporters of a particular instance of the Graded Incoherence Principle should also endorse the corresponding instance of the Accuracy-Dominance Principle. Additionally, if promoting accuracy reduces incoherence, then the corresponding instance of the Accuracy-Dominance Principle is compatible with the instance of the Graded Incoherence Principle, for it could never happen that one norm says a credence function c is epistemically better than c' while the other says the opposite. Consider the following two schemata:

> (IIIa) *Reducing inaccuracy promotes coherence* (*weak dominance version*): For all credence functions c, c' defined on the same set of propositions, if c weakly $\mathfrak{I}$-accuracy-dominates c', then c is not more $\mathcal{I}_d$-incoherent than c'.

> (IIIb) *Reducing inaccuracy promotes coherence* (*strong dominance version*): For all credence functions c, c' defined on the same set of propositions, if c strongly $\mathfrak{I}$-accuracy-dominates c', then c is less $\mathcal{I}_d$-incoherent than c'.

As we show in the appendix, every additive inaccuracy measure $\mathfrak{I}$ that is based on a continuous, strictly proper scoring rule defines a corresponding divergence $d_{\mathfrak{I}}$ that can be used to measure incoherence with $\mathcal{I}_{d_{\mathfrak{I}}}$. For example, if we use the Brier score as our inaccuracy measure, the corresponding incoherence measure is the Euclidean distance (or squared Euclidean distance) to some closest coherent credence function. When inaccuracy is measured with the additive logarithmic scoring rule, the matching divergence

https://doi.org/10.1086/690715 Published online by Cambridge University Press

to measure incoherence is a variation of the Kullback-Leibler divergence (Capotorti, Regoli, and Vattari 2009). For more on the relation between scoring rules and divergences, see Gneiting and Raftery (2007).

By matching each inaccuracy measure with its corresponding divergence, we can generate true instances of IIIa and IIIb for every additive inaccuracy measure $\mathfrak{I}$ that is based on a continuous, strictly proper scoring rule:

> **Theorem 1.** Let $c : F \to [0, 1]$ and $c' : F \to [0, 1]$ be credence functions, and $\mathfrak{I}$, an inaccuracy measure generated by a strictly proper scoring rule. If c' weakly $\mathfrak{I}$-accuracy-dominates c, then $\mathcal{I}_{d_\mathfrak{I}}(c') \leq \mathcal{I}_{d_\mathfrak{I}}(c)$; if c' strongly $\mathfrak{I}$-accuracy-dominates c, then $\mathcal{I}_{d_\mathfrak{I}}(c') < \mathcal{I}_{d_\mathfrak{I}}(c)$.

Intuitively, the divergence $d_\mathfrak{I}(c, g)$, based on the inaccuracy measure $\mathfrak{I}$, measures the expected value (according to g) of $\mathfrak{I}$-inaccuracy that one gains when using c instead of g. Hence, when c strongly (weakly) accuracy dominates c', the expected value of the $\mathfrak{I}$-inaccuracy of c, according to any coherent g, is less than (less than or equal to) the expected value of the $\mathfrak{I}$ -inaccuracy of c', for the same g. That is, $d_\mathfrak{I}(c, g) < d_\mathfrak{I}(c', g)$ (or $d_\mathfrak{I}(c, g) \leq d_\mathfrak{I}(c', g)$). By the definition of $\mathcal{I}_{d_\mathfrak{I}}$, the result above follows.

This is not to say that the only true instances of IIIa and IIIb are the ones that are generated by pairing inaccuracy measures with their associated incoherence measures. We might be able to "mix and match" inaccuracy measures and incoherence measures in some cases. Finding a general result about the extent to which this is possible is beyond the scope of this article. However, we show in the appendix that, if the Brier score is chosen as our inaccuracy measure, and we want to select an incoherence measure that is based on a p-norm distance measure, the only incoherence measure for which IIIa and IIIb hold is (squared) Euclidean distance to some closest coherent credence function (proposition 3).

We have thus established our first central result: for suitable combinations of $\mathcal{I}_d$ and $\mathfrak{I}$, one can never achieve a certain gain in accuracy by becoming more incoherent.[10] There can thus never be a conflict between the goals of coherence and accuracy when incoherence and inaccuracy are measured with appropriate combinations of $\mathcal{I}_d$ and $\mathfrak{I}$, in the sense that accuracy considerations might recommend that the agent should become more incoherent or vice versa. Moreover, for suitable combinations of $\mathcal{I}_d$ and $\mathfrak{I}$, if the

10. By "a certain gain in accuracy," we mean a situation in which at least a relation of weak dominance holds between the credence function the agent moves to and the agent's initial credence function.

https://doi.org/10.1086/690715 Published online by Cambridge University Press

agent moves to a credence function that is more accurate than her current credences in every possible world, then her credences will always become less incoherent. In other words, if one endorses the Graded Incoherence Principle, one should also endorse the Strong Accuracy-Dominance Principle for suitable combinations of $\mathcal{I}_d$ and $\mathfrak{I}$.

We now return to claims IIa* and IIb* in order to see whether a weakened version of them might be true. The claims we have already proved—IIIa and IIIb—can only encourage a proponent of the Graded Incoherence Principle to endorse an Accuracy-Dominance Principle, since it promotes coherence, but not the other way round. A weaker version of IIa* or IIb* could state that at least some way of reducing incoherence promotes accuracy. And there is in fact a specific way of reducing incoherence in one's credences that always constitutes an improvement of one's epistemic situation. Assuming that we choose a convex inaccuracy measure, and its associated incoherence measure, we can show that an agent is guaranteed to improve her accuracy if she moves to a less incoherent credence function that lies on the direct path between her initial credence function and some closest coherent credence function. Informally speaking, an accuracy measure is convex just in case the reward for moving one's credence toward the truth is not greater than the penalty for moving one's credence away from the truth by the same amount.[11]

> **Proposition 1.** Let $c : F \to [0, 1]$ be an incoherent credence function, $\mathfrak{I}$ be a convex inaccuracy measure generated by a continuous, strictly proper scoring rule, and $c^* : F \to [0, 1]$ be a coherent credence function minimizing $d_{\mathfrak{I}}(c, c^*)$. For any λ within the interval (0,1], let c_λ be a credence function such that $c_\lambda(\varphi) = (1 - \lambda)c(\varphi) + \lambda c^*(\varphi)$ for all $\varphi \in F$; c_λ strongly $\mathfrak{I}$-accuracy-dominates c for any $\lambda \in (0, 1]$.

Having proved this proposition, we now know how to weaken IIb*, so that we can generate true instances of it by substituting matching pairs of inaccuracy and incoherence measures $\mathfrak{I}$, $\mathcal{I}_d$:

11. Convexity is considered to be an attractive feature of scoring rules and their associated inaccuracy measures because it encourages agents to be conservative in revising their beliefs. Suppose we could choose to undergo a random process that is equally likely to raise or lower our credences slightly. In this scenario, a nonconvex inaccuracy measure might deem this process to be on average beneficial for the accuracy of our credences. But many philosophers think that one should not revise one's beliefs on the basis of such a random process. Requiring that scoring rules, and hence inaccuracy measures, must be convex is sufficient to discourage such a belief revision strategy (Joyce 2009).

https://doi.org/10.1086/690715 Published online by Cambridge University Press

(II′) *Approximating some closest coherent credence function promotes accuracy*: If c is an incoherent credence function on F, and c^* is some d-closest coherent credence function to c, then c^* strongly $\mathfrak{I}$-accuracy-dominates c. Moreover, for any $\lambda \in (0, 1)$, let c_λ be a credence function on F such that $c_\lambda(\varphi) = (1 - \lambda)\, c(\varphi) + \lambda c^*(\varphi)$ for all $\varphi \in F$; then c_λ strongly $\mathfrak{I}$-accuracy-dominates c for any $\lambda \in (0, 1)$.

We can now state our second central result: assume that inaccuracy is measured by a convex inaccuracy measure based on a continuous, strictly proper scoring rule and incoherence is measured on the basis of the divergence that corresponds to the inaccuracy measure. Any such combination of inaccuracy and incoherence measures guarantees that if the agent reduces the incoherence of her credence function by adopting a credence function that lies on the direct path between her initial credence function and some closest coherent credence function, then her new credence function will be more accurate in every possible world than her initial credence function. It is a question for further research whether the result can be generalized to nonconvex inaccuracy measures.

On the basis of our results regarding schema II′, we can now state a weaker version of the Graded Incoherence Principle:

Weak Graded Incoherence Principle: For any credence functions $c : F \to [0, 1]$ and $c' : F \to [0, 1]$, if c is less $\mathcal{I}_d$-incoherent than c' and lies on the direct path between c' and some d-closest coherent credence function, then c is epistemically better than c'.

Our results show that the Weak Graded Incoherence Principle must be endorsed by any accuracy-firster who employs a convex, additive inaccuracy measure based on a continuous, strictly proper scoring and a suitably associated incoherence measure. Moreover, we include some preliminary results in the appendix that show that there are limitations on "mixing and matching" inaccuracy measures and incoherence measures in this context as well. According to proposition 4, the first part of II′ does not hold if we choose the Brier score as our inaccuracy measure and combine it with any incoherence measure based on a p-norm distance other than $p = 2$.[12]

12. As shown in the appendix, this result relies on a proof in which incoherence results from assigning different credences to logically equivalent propositions. This kind of example precludes identifying propositions with sets of possible worlds. Since this seems like a plausible way in which an agent could be incoherent, we do not consider this to be a problem.

https://doi.org/10.1086/690715 Published online by Cambridge University Press

5. Conclusion and Prospects. We have established two interesting results for the accuracy-firster: first, she must endorse qualified instances of the claim that being less incoherent is epistemically better, which we have schematically represented in the Weak Graded Incoherence Principle. Second, due to the compatibility result IIIa, none of the relevant instances of the Accuracy-Dominance Principles favored by the accuracy-firster entail results that are incompatible with the Graded Incoherence Principle (assuming we choose "matching" $\mathcal{I}_d$ and $\mathfrak{I}$). Moreover, we noticed along the way that, because of claim IIIb, a proponent of the Graded Incoherence Principle should endorse the Strong Accuracy-Dominance Principle since it promotes coherence (again, assuming we choose instances of the principles based on "matching" $\mathcal{I}_d$ and $\mathfrak{I}$).

We also showed that whether instances of the principle schemata are true depends on the chosen combinations of inaccuracy measures and incoherence measures. Each inaccuracy measure has a naturally associated incoherence measure, and we showed some preliminary results about the Brier score not being combinable with incoherence measures based on p-norm distances with $p \neq 2$ to get true instances of the principle schemata II′, IIIa, and IIIb. Hence, considering only p-norm distance measures, a proponent of accuracy-first epistemology who considers the Brier score to be the appropriate measure of inaccuracy, and who wants to avoid contradicting the Graded Incoherence Principle, is committed to claiming that the incoherence of a credence function must be its (squared) Euclidean distance to some closest coherent credence function. An interesting question for future work is whether a version of the Graded Incoherence Principle can be established for accuracy-firsters who endorse inaccuracy measures that are not convex. Moreover, there is more work to be done in showing which inaccuracy measures and incoherence measures can be combined.

Our results give us a novel way of evaluating the merits of different incoherence and inaccuracy measures. (In)compatibility results regarding which combinations of measures $\mathcal{I}_d$ and $\mathfrak{I}$ produce true and false instances of the principle schemata II′, IIIa, and IIIb provide us with a license to translate arguments for or against a particular measure to its associated incoherence or inaccuracy measure. For example, one could argue against measuring incoherence by determining the absolute distance to some closest coherent credence function by arguing for the Brier score as the correct inaccuracy measure and pointing out that the Brier score combined with the absolute incoherence measure forces us to reject the relevant instances of the principle schemata II′, IIIa, and IIIb. Or one could argue against a particular inaccuracy measure by arguing for an incompatible incoherence measure and the relevant instances of the principle schemata. Of course, only someone who is interested in upholding the relevant instances of II′, IIIa, and IIIb would find this argumentative strategy attractive.

https://doi.org/10.1086/690715 Published online by Cambridge University Press

Appendix

Given a credence function $c:\{\varphi_1, \ldots, \varphi_m\} \to [0,1]$ such that $c(\varphi_i) = q_i$ for $1 \leq i \leq m$, we say that $q = \langle q_1, \ldots, q_m \rangle$ is its corresponding vector. For a given possible world w_i, recall that $I_{w_i} : \mathcal{L} \to \{0, 1\}$ is its indicator function. Abusing the notation, we say that $v^i = \langle I_{w_i}(\varphi_1), \ldots, I_{w_i}(\varphi_m) \rangle \in \{0, 1\}^m$ is its corresponding vector.

Suppose that regarding the set of propositions $F = \{\varphi_1, \ldots, \varphi_m\}$ there are only ℓ distinct possible worlds. Let $V = \{v^1, \ldots, v^\ell\} \subset \{0, 1\}^m$ be the set of their corresponding vectors. A credence function $c : F \to [0, 1]$ is coherent iff its corresponding vector is in the convex hull of V, which we denote by $\mathrm{CH}(V) \subset \mathbb{R}^m$ (De Finetti 1974).

When an inaccuracy measure $\mathfrak{I}$ is generated by a scoring rule s, for arbitrary $q \in \mathrm{CH}(V)$ and $q' \in [0, 1]^m$, $\mathfrak{I}(q', q)$ denotes the expected value of $\mathfrak{I}(q', v)$ when $v \in V$ is distributed according to any probability measure π satisfying the coherent credence function c corresponding to q:

$$\mathfrak{I}(q', q) = \sum_{j=1}^{\ell} \pi(w_j)\mathfrak{I}(q', v^j).$$

With some algebraic manipulation, any π satisfying c leads to

$$\mathfrak{I}(q', q) = \sum_{i=1}^{m} (q_i s(q'_i, 1) + (1 - q_i)s(q'_i, 0)).$$

For any q corresponding to a possible world $v \in V$, the definitions of $\mathfrak{I}$ collapse; $\mathfrak{I}(q', q) = \mathfrak{I}(q', v)$.

Every additive inaccuracy measure $\mathfrak{I} : [0, 1]^m \times V \to \mathbb{R}$ based on a strictly proper scoring rule defines a divergence $d_{\mathfrak{I}} : [0, 1]^m \times \mathrm{CH}(V) \to [0, \infty)$ in the following way:

$$d_{\mathfrak{I}}(q', q) = \mathfrak{I}(q', q) - \mathfrak{I}(q, q).$$

In other words, $d_{\mathfrak{I}}(c', c)$ measures the expected gain in inaccuracy, according to a coherent c, due to adopting c' instead of c.

As $\mathfrak{I}$ is generated by a strictly proper scoring rule s, $\mathfrak{I}(q, q) \leq \mathfrak{I}(q', q)$ and $d_{\mathfrak{I}}(q', q) \geq 0$; furthermore, $d_{\mathfrak{I}}(q', q) = 0$ iff $q = q'$.

We say an inaccuracy measure $\mathfrak{I} : [0, 1]^m \times V \to \mathbb{R}$ is convex if, for every $\langle q, q', v, \lambda \rangle \in [0, 1]^m \times [0, 1]^m \times V \times [0, 1]$, we have

$$\mathfrak{I}(\lambda q + (1 - \lambda)q', v) \leq \lambda\mathfrak{I}(q, v) + (1 - \lambda)\mathfrak{I}(q', v).$$

The proposition below shows that claims IIa and IIb have no true instances for inaccuracy measures generated by continuous, strictly proper scoring rules:

Proposition 2. For every inaccuracy measure $\mathfrak{I}$ generated by a continuous, strictly proper scoring rule and for any divergence d, there are credence func-

https://doi.org/10.1086/690715 Published online by Cambridge University Press

tions $c : F \to [0, 1]$ and $c' : F \to [0, 1]$ such that $\mathcal{I}_d(c) < \mathcal{I}_d(c')$ but c does not weakly $\mathfrak{I}$-accuracy-dominate c'.

Proof. Let F be $\{\varphi, \neg\varphi\}$ for a consistent, nontautological proposition $\varphi \in \mathcal{L}$. There are two distinct possible worlds, w_1 and w_2, corresponding to the vectors $v^1 = \langle 1, 0 \rangle$ and $v^2 = \langle 0, 1 \rangle$, respectively. Let $c : F \to [0, 1]$ and $c' : F \to [0, 1]$ be such that $c(\varphi) = 1$ and $c(\neg\varphi) = 0$, corresponding to the vector $q = \langle 1, 0 \rangle = v^1$, and $c'(\varphi) = 0$ and $c'(\neg\varphi) = 1$, corresponding to the vector $q' = \langle 0, 1 \rangle = v^2$. As both are coherent and $\mathcal{I}_d$ satisfies consistency, $\mathcal{I}_d(c) = \mathcal{I}_d(c') = 0$. Furthermore, since $\mathfrak{I}$ is generated by a strictly proper scoring rule s, c is the only credence function whose corresponding vector q minimizes $\mathfrak{I}(q, v^1)$. Hence, $\mathfrak{I}(q, v^1) < \mathfrak{I}(q', v^1)$, and, by the density of the real numbers, there is a $k \in \mathbb{R}$ such that $\mathfrak{I}(q, v^1) < k < \mathfrak{I}(q', v^1)$.

Now consider a credence function $c_\varepsilon : F \to [0, 1]$ such that $c_\varepsilon(\varphi) = 1$ and $c_\varepsilon(\neg\varphi) = \varepsilon$, corresponding to the vector $q_\varepsilon = \langle 1, \varepsilon \rangle$, for some $\varepsilon \in (0, 1]$. The inaccuracy measure $\mathfrak{I}$ is the sum of continuous scoring rules, so it is continuous itself. Thus,

$$\lim_{\varepsilon \to 0^+} \mathfrak{I}(q_\varepsilon, v^1) = \mathfrak{I}(q, v^1) < k < \mathfrak{I}(q', v^1).$$

By the limit definition, there is a δ such that $0 < \varepsilon < \delta$ implies $\mathfrak{I}(q_\varepsilon, v^1) \leq k < \mathfrak{I}(q', v^1)$. Finally, taking such an ε in (0, 1), we have an incoherent credence function c_ε, with $\mathcal{I}_d(c_\varepsilon) > 0 = \mathcal{I}_d(c')$, which is not even weakly $\mathfrak{I}$-accuracy-dominated by the coherent c'. QED

Proof of Theorem 1. Suppose that $F = \{\varphi_1, \ldots, \varphi_m\}$ and that there are exactly ℓ distinct possible worlds, corresponding to the vectors in the set $V = \{v^1, \ldots, v^\ell\}$. Let $q, q' \in \mathbb{R}^m$ be the corresponding vectors of c and c', respectively. Let $g = \langle g_1, \ldots, g_m \rangle$ and g' be vectors in $\mathrm{CH}(V)$ such that $\mathcal{I}_\mathfrak{I}(c) = d(q, g)$ and $\mathcal{I}_\mathfrak{I}(c') = d(q', g')$. The distance from q' to a closest vector $g' \in \mathrm{CH}(V)$ cannot be greater than the distance from q' to an arbitrary vector in $\mathrm{CH}(V)$. Therefore, $\mathcal{I}_{d_\mathfrak{I}}(c') = d_\mathfrak{I}(q', g') \leq d_\mathfrak{I}(q', g)$; thus, $d_\mathfrak{I}(q', g) \leq d_\mathfrak{I}(q, g)$ would imply $\mathcal{I}_{d_\mathfrak{I}}(c') \leq \mathcal{I}_{d_\mathfrak{I}}(c)$. Similarly, $d_\mathfrak{I}(q', g) < d_\mathfrak{I}(q, g)$ entails $\mathcal{I}_{d_\mathfrak{I}}(c') < \mathcal{I}_{d_\mathfrak{I}}(c)$. We will prove that if c' weakly (strongly) $\mathfrak{I}$-accuracy-dominates c, then $d_\mathfrak{I}(q', g) \leq d_\mathfrak{I}(q, g)$ ($d_\mathfrak{I}(q', g) < d_\mathfrak{I}(q, g)$).

By the definition of $\mathfrak{I}$, we have, for a π satisfying g:

$$\mathfrak{I}(q, g) = \sum_{j=1}^{\ell} \pi(w_j) \mathfrak{I}(q, v^j); \quad \mathfrak{I}(q', g) = \sum_{j=1}^{\ell} \pi(w_j) \mathfrak{I}(q', v^j).$$

If c' strongly $\mathfrak{I}$-accuracy-dominates c, $\mathfrak{I}(q', v^j) < \mathfrak{I}(q, v^j)$ for all $v^j \in V$. As $\sum_{j=1}^{\ell} \pi(w_j) = 1$, there must be a possible world w_j such that $\pi(w_j) > 0$. Furthermore, $\pi(w_j) \geq 0$; thus, $\mathfrak{I}(q', g) < \mathfrak{I}(q, g)$. Subtracting $\mathfrak{I}(g, g)$ from both

https://doi.org/10.1086/690715 Published online by Cambridge University Press

sides, we obtain $d_{\mathfrak{I}}(q', g) < d_{\mathfrak{I}}(q, g)$. If the accuracy dominance is weak, $\mathfrak{I}(q', v^j) \leq \mathfrak{I}(q, v^j)$ for all $v^j \in V$, $\mathfrak{I}(q', g) \leq \mathfrak{I}(q, g)$, and $d_{\mathfrak{I}}(q', g) \leq d_{\mathfrak{I}}(q, g)$. QED

Note that if we replace $\mathcal{I}_{d_{\mathfrak{I}}}$ by $f \circ \mathcal{I}_{d_{\mathfrak{I}}}$ for any monotone function f, the result proved above still holds. For instance, when $\mathfrak{I}$ is the Brier score, $d_{\mathfrak{I}}$ is the squared Euclidean distance. But if incoherence is measured using $\mathcal{I}_d$ and d is Euclidean distance, the result above still applies.

We can show by counterexamples that this holds among the p-norm distances only for $p = 2$. That is, if incoherence is measured using a p-norm distance (to some closest coherent credences) other than Euclidean distance, then incoherence can increase when accuracy (measured with the Brier score) increases in all worlds. We denote the incoherence measure $\mathcal{I}_d$ when d is some p-norm distance simply by $\mathcal{I}_p$.

Proposition 3. For any $p \in \mathbb{N}_{>0} \cup \{\infty\}$ with $p \neq 2$, there are credence functions $c : F \to [0, 1]$ and $c' : F \to [0, 1]$ such that c' strongly Brier-accuracy-dominates c, but $\mathcal{I}_p(c') > \mathcal{I}_p(c)$.

Proof. We give two counterexamples: the first for $p = 1$ and the second for all $p \geq 3$. For the first counterexample, let $F = \{\varphi_1, \varphi_2, \varphi_3\}$ form a partition. Let $c : F \to [0, 1]$ and $c' : F \to [0, 1]$ be credence functions such that $c(\varphi_1) = c(\varphi_2) = 0.7$, $c(\varphi_3) = 0.05$ and $c'(\varphi_1) = c'(\varphi_2) = 0.68$ $c'(\varphi_3) = 0.11$. Regarding F, there are three distinct possible wolds w_1, w_2, w_3, such that φ_i is true in w_j iff $i = j$. Computing the Brier scores in each possible world, we can see that c' strongly accuracy dominates c:

$$\langle B(c, w_1), B(c, w_2), B(c, w_3) \rangle = \langle 0.5825, 0.5825, 1.8825 \rangle.$$
$$\langle B(c', w_1), B(c', w_2), B(c', w_3) \rangle = \langle 0.5769, 0.5769, 1.7169 \rangle.$$

Considering $q = \langle 0.7, 0.7, 0.05 \rangle$ and $q' = \langle 0.68, 0.68, 0.11 \rangle$, note that $\Sigma_i q_i = 1.45$; thus, for c to be coherent, the absolute changes in q must sum to at least 0.45. Consequently, $\mathcal{I}_1(c) \geq 0.45$; similarly, $\mathcal{I}_1(c') \geq 0.47$. Defining $g = \langle 0.5, 0.5, 0 \rangle$, we also have that $\mathcal{I}_1(c) \leq \| q - g \|_1 = 0.45$ and $\mathcal{I}_1(c') \leq \| q' - g \| \mathcal{I}_1 = 0.47$. Hence, $\mathcal{I}_1(c) = 0.45 < 0.47 = \mathcal{I}_1(c')$, finishing the first part of the proof.

Now consider a partition $F = \{\varphi_1, \varphi_2, \ldots, \varphi_{1,002}\}$. Let $c : F \to [0, 1]$ and $c' : F \to [0, 1]$ be credence functions such that $c(\varphi_1) = c(\varphi_2) = 0.6$, $c(\varphi_3) = \ldots = c(\varphi_{1,002}) = 0.0012$ and $c'(\varphi_1) = c'(\varphi_2) = 0.60001$, $c'(\varphi_3) = \ldots = c'(\varphi_{1,002}) = 0.0011$. Again, let w_i be the possible world in which only φ_i is true. Computing the Brier scores in each possible world, we can see that c' strongly accuracy dominates c—note that any world w_i with $i \geq 3$ yields the same Brier scores for c and c':

https://doi.org/10.1086/690715 Published online by Cambridge University Press

$$\langle B(c, w_1), B(c, w_2), B(c, w_{i\geq 3}) \rangle \cong \langle 0.52144, 0.52144, 1.71904 \rangle.$$
$$\langle B(c', w_1), B(c', w_2), B(c', w_{i\geq 3}) \rangle \cong \langle 0.52121, 0.52121, 1.71903 \rangle.$$

Let q, $q' \in \mathbb{R}^{1,002}$ be the vectors corresponding to c and c', respectively. We leave for the reader to note that $g^* = \langle 0.5, 0.5, 0, \ldots, 0 \rangle \in \mathbb{R}^{1,002}$ is such that $\mathcal{I}_p(c) = \| q - g^* \|_p$ and $\mathcal{I}_p(c') = \| q' - g^* \|_p$ for any $p \geq 3$; that is, g^* corresponds to a closest coherent credence function for both c and c' regarding any p-norm distance with some $p \geq 3$.

Intuitively, the closest $g = \langle g_1, \ldots, g_{1,002} \rangle$ has $g_1 = g_2$ and $g_3 = \ldots = g_{1,002}$, so that $g_1 = g_2 = (1 - 1,000g_3)/2$. One can prove that any $g_3 > 0$ implies $\| q - g^* \|_p < \| q - g \|_p$ and $\| q' - g^* \|_p < \| q' - g \|_p$ for all $p \geq 3$.

Back to measuring $\mathcal{I}_p(c)$ and $\mathcal{I}_p(c')$, the case $p = \infty$ is straightforward:

$$\mathcal{I}_\infty(c') = \max_i \left| q'_i - g^*_i \right| = 0.10001 > 0.1 = \max_i \left| q_i - g^*_i \right| = \mathcal{I}_\infty(c).$$

For any integer $p \geq 3$, we can write $\mathcal{I}_p(c)$ and $\mathcal{I}_p(c')$ using q, q', and g^*:

$$\begin{aligned} \mathcal{I}_p(c) &= \sqrt[p]{2(0.6 - 0.5)^p + 1,000(0.0012)^p} \\ &= \sqrt[p]{2 \times 0.1^p + 1,000 \times 0.0012^p}. \\ \mathcal{I}_p(c') &= \sqrt[p]{2(0.60001 - 0.5)^p + 1,000(0.0011)^p} \\ &= \sqrt[p]{2 \times 0.10001^p + 1,000 \times 0.0011^p}. \end{aligned}$$

As $\mathcal{I}_p(c') > \mathcal{I}_p(c)$ iff $\mathcal{I}_p(c')^p - \mathcal{I}_p(c)^p > 0$, we define a function $f : \mathbb{N}_{>0} \to \mathbb{R}$:

$$\begin{aligned} f(p) &= \mathcal{I}_p(c')^p - \mathcal{I}_p(c)^p \\ &= 2(0.10001^p - 0.1^p) + 1,000(0.0011^p - 0.0012^p). \end{aligned}$$

We use induction to prove that $f(p) > 0$ for all $p \geq 3$. Note that $f(3) \cong 1.0001 > 1 > 0$, which is the base of the induction. Now we compute $f(p + 1)$:

$$\begin{aligned} f(p + 1) &= 2(0.10001 \times 0.10001^p - 0.1 \times 0.1^p) \\ &\quad + 1{,}000(0.0011 \times 0.0011^p - 0.0012 \times 0.0012^p). \\ f(p + 1) &= 10^{-5} 0.10001^p 2 - 0.1 \times 0.0012^p \\ &\quad + 2 \times 0.1(0.10001^p - 0.1^p) \\ &\quad + 1{,}000 \times 0.0011(0.0011^p - 0.0012^p). \end{aligned} \tag{A1}$$

Using the induction hypothesis, as $0.1 > 0.0011$, we obtain

https://doi.org/10.1086/690715 Published online by Cambridge University Press

$$f(p) = 2(0.10001^p - 0.1^p) + 1{,}000(0.0011^p - 0.0012^p) > 0.$$
$$0.1 \times 2(0.10001^p - 0.1^p) > 0.0011 \times 1{,}000(0.0012^p - 0.0011^p). \tag{A2}$$

But, for any $p \geq 3 > \ln 5{,}000/(\ln 0.10001 - \ln 0.0012)$, we have

$$10^{-5}0.10001^p 2 - 0.1 \times 0.0012^p > 0. \tag{A3}$$

Applying (A2) and (A3) in (A1) we prove that $f(p + 1) > 0$ for all $p \geq 3$. Thus, $\mathcal{I}_p(c') > \mathcal{I}_p(c)$ for all $p \geq 3$, finishing the proof. QED

Proof of Proposition 1. Predd et al. (2009) proved the case $\lambda = 1$, together with the uniqueness of the closest coherent credence function c^*. Hence, taking those results for granted, we focus on $\lambda \in (0, 1)$.

Consider a c_λ for an arbitrary $\lambda \in (0, 1)$. Due to the convexity of $\mathfrak{I}$, for every possible world $v \in V$ we have

$$\mathfrak{I}(c_\lambda, v) \leq \lambda\mathfrak{I}(c, v) + (1 - \lambda)\mathfrak{I}(c^*, v). \tag{A4}$$

For every vector $v \in V$, applying the result of Predd et al. (2009) that $\mathfrak{I}(c^*, v) < \mathfrak{I}(c, v)$, we obtain

$$\mathfrak{I}(c_\lambda, v) \leq \lambda\mathfrak{I}(c, v) + (1 - \lambda)\mathfrak{I}(c^*, v) < \lambda\mathfrak{I}(c, v) + (1 - \lambda)\mathfrak{I}(c, v).$$
$$\mathfrak{I}(c_\lambda, v) < \mathfrak{I}(c, v).$$

Finally, c_λ strongly $\mathfrak{I}$-accuracy-dominates c. QED

We show below by counterexamples that when inaccuracy is measured via the Brier score, and incoherence via any p-norm distance with $p \neq 2$, the first part of claim II′ does not hold. For a credence function $c : F \to [0, 1]$ corresponding to a vector q, we say that $c^* : F \to [0, 1]$ is the p-closest coherent credence function if its corresponding vector q^* uniquely minimizes $\| q - q^* \|_p$.

Proposition 4. There is a credence function c such that, for any $p \in \mathbb{N}_{>0} \cup \{\infty\}$ with $p \neq 2$, the p-closest coherent credence function c^* does not weakly accuracy dominate c according to Brier score.

Proof. Consider the set of propositions $F = \{\varphi, \varphi', \neg\varphi\}$, where φ is logically equivalent to φ'. Define the credence function $c : F \to [0, 1]$ with $c(\varphi) = c(\varphi') = c(\neg\varphi) = 0.6$ and define $q = \langle 0.6, 0.6, 0.6 \rangle$. These propositions yield a set V of two vectors corresponding to the possible worlds: w_1, where φ is true, and w_2, where φ is false. For $p = 1$, $g^* = \langle 0.6, 0.6, 0.4 \rangle$ is the unique value of $g \in \mathrm{CH}(V)$ that minimizes $\| q - g \|_1$; g^* corresponds to the credence function c^* with $c^*(\varphi) = c^*(\varphi') = 0.6$ and $c^*(\neg\varphi) = 0.4$. To

https://doi.org/10.1086/690715 Published online by Cambridge University Press

see this, consider the alternatives. If $g_3 = 0.4 + \varepsilon > 0.4$, for some $\varepsilon \in (0, 0.6]$, then coherence implies $g_1 = g_2 = 0.6 - \varepsilon$ and $\| q - g \|_1 \geq 0.2 + \varepsilon > 0.2 = \| q - g^* \|_1$; that is, g^* is closer to q than g is. Similarly, if $g_3 = 0.4 - \varepsilon < 0.4$ for some $\varepsilon \in (0, 0.4]$, coherence implies $g_1 = g_2 = 0.6 + \varepsilon$ and $\| q - g \|_1 = 0.2 + 3\varepsilon > 0.2 = \| q - g^* \|_1$. Hence, $\mathcal{I}_1(c) = \| q - g^* \|_1 = 0.2$. Computing the Brier scores,

$$\langle B(c, w_1), B(c, w_2) \rangle = \langle 0.68, 0.88 \rangle;$$
$$\langle B(c^*, w_1), B(c^*, w_2) \rangle = \langle 0.48, 1.08 \rangle.$$

Therefore, c^* does not weakly accuracy dominate c. For $p = \infty$, the vector corresponding to the closest coherent credence function c^* is $g^* = \langle 0.5, 0.5, 0.5 \rangle$, yielding $\mathcal{I}_\infty(c) = \| \langle 0.6, 0.6, 0.6 \rangle - \langle 0.5, 0.5, 0.5 \rangle \|_\infty = 0.1$. If $g_3 = 0.5 + \varepsilon$, for some $\varepsilon \in (0, 0.5]$, coherence implies $g_1 = g_2 = 0.5 - \varepsilon$ and $\| p - g \|_\infty = 0.1 + \varepsilon > \| p - g^* \|_\infty$. If $g_3 < 0.5$, $\| q - g \|_\infty = 0.6 - g_3 > 0.1 = \| q - g^* \|_\infty$. The Brier scores are $B(c^*, w_1) = B(c^*, w_2) = 0.75$, and no accuracy-dominance relation holds between c and c^*.

Now we consider all integers $p \geq 3$. First, note that any vector corresponding to a coherent credence function must have the form $\langle x, x, 1 - x \rangle$ for some $x \in [0, 1]$. Furthermore, x must lie within [0.4, 0.6]. If $x = g_1 = g_2 > 0.6$, then $g_3 < 0.4$, and the vector $g' = \langle 0.6, 0.6, 0.4 \rangle$ would be (p-) closer to q than g would for any p since $|g'_i - q_i| < |g_i - q_i|$ for all $1 \leq i \leq 3$. Similarly, if $x = g_1 = g_2 < 0.4$, $g' = \langle 0.4, 0.4, 0.6 \rangle$ would be closer to q. Note that the vector $g = \langle x, x, 1 - x \rangle$ minimizes $\| q - g \|_p$ iff it minimizes $\| q - g \|_p^p$. Hence, for a given p, we look for the $x \in [0.4, 0.6]$ that minimizes the function h_p: $[0.4, 0.6] \to \mathbb{R}$, defined as

$$h_p(x) = |x - 0.6|^p + |x - 0.6|^p + |1 - x - 0.6|^p = 2(0.6 - x)^p + (x - 0.4)^p.$$

We take the derivative $h'_p(x)$ with respect to x and make it equal to 0 to obtain the critical points:

$$-2p(0.6 - x)^{p-1} + p(x - 0.4)^{p-1} = 0; \quad x = \frac{0.6\sqrt[p-1]{2} + 0.4}{\sqrt[p-1]{2} + 1}. \tag{A5}$$

For a given $p \geq 3$, define x_p as the value given by (A5). The reader can note that $0.6 \geq x_p \geq 0.4$ for any $p \geq 3$, for x_p is a weighted average between 0.6 and 0.4—with weights $\sqrt[p-1]{2}$ and 1. For any $p \geq 3$, as $h''_p(x) > 0$ for all $x \in [0.4, 0.6]$, x_p is the minimum of $h_p(x)$, and $g_p = \langle x_p, x_p, 1 - x_p \rangle$ corresponds to the p-closest coherent credence function c_p^*. For an arbitrary $p \geq 3$, we compute the Brier score of the p-closest coherent function c_p^* in w_1:

$$B(c_p^*, w_1) = \| \langle 1, 1, 0 \rangle - g_p \|_2^2 = 2(1 - x_p)^2 + (x_p - 1)^2$$
$$= 3(1 - x_p)^2 = f(p).$$

https://doi.org/10.1086/690715 Published online by Cambridge University Press

One can compute $B(c_3^*, w_1) = f(3) \cong 0.699 > 0.68 = B(c, w_1)$ to see that c_p^* does not weakly accuracy dominate c for $p = 3$. Note that $f(p)$ increases when x_p decreases as long as $x_p \leq 1$. As a weighted average, x_p gets closer to 0.4 if the weight associated to 0.6 ($\sqrt[p-1]{2}$) decreases. Finally, for any $p \geq 4$, $\sqrt[p-1]{2} < \sqrt[3-1]{2}$ and $x_p < x_3$; therefore, $B(c_p^*, w_1) = f(p) > f(3) = B(c_3^*, w_1) > B(c, w_1)$, and c^* does not weakly accuracy dominate c. QED

REFERENCES

Caie, M. 2013. "Rational Probabilistic Incoherence." *Philosophical Review* 122 (4): 527–75.

Capotorti, A., G. Regoli, and F. Vattari. 2009. "On the Use of a New Discrepancy Measure to Correct Incoherent Assessments and to Aggregate Conflicting Opinions Based on Imprecise Conditional Probabilities." In *Proceedings of the Sixth International Symposium on Imprecise Probability: Theories and Applications, ISIPTA*, ed. Thomas Augustin. [England]: Society for Imprecise Probability.

Christensen, D. 2007. "Does Murphy's Law Apply in Epistemology? Self-Doubt and Rational Ideals." *Oxford Studies in Epistemology* 2:3–31.

De Bona, G., and M. Finger. 2015. "Measuring Inconsistency in Probabilistic Logic: Rationality Postulates and Dutch Book Interpretation." *Artificial Intelligence* 227:140–64.

De Finetti, B. 1974. *Theory of Probability*. Chichester: Wiley.

Easwaran, K. 2013. "Expected Accuracy Supports Conditionalization—and Conglomerability and Reflection." *Philosophy of Science* 80 (1): 119–42.

Gneiting, T., and A. E. Raftery. 2007. "Strictly Proper Scoring Rules, Prediction, and Estimation." *Journal of the American Statistical Association* 102 (477): 359–78.

Goldman, A. 2002. *Pathways to Knowledge: Public and Private*. Oxford: Oxford University Press.

Greaves, H., and D. Wallace. 2006. "Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility." *Mind* 115 (459): 607–32.

Hacking, I. 1967. "Slightly More Realistic Personal Probability." *Philosophy of Science* 34:311–25.

Harman, G. 1986. *Change in View*. Cambridge, MA: MIT Press.

Joyce, J. M. 1998. "A Nonpragmatic Vindication of Probabilism." *Philosophy of Science* 65 (4): 575–603.

———. 2009. "Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief." In *Degrees of Belief*, ed. F. Huber and C. Schmidt-Petri, 263–97. Synthese Library 342. Dordrecht: Springer.

Moss, S. 2011. "Scoring Rules and Epistemic Compromise." *Mind* 120 (480): 1053–69.

Pettigrew, R. 2016. *Accuracy and the Laws of Credence*. Oxford: Oxford University Press.

Predd, J. B., R. Seiringer, E. H. Lieb, D. N. Osherson, H. V. Poor, and S. R. Kulkarni. 2009. "Probabilistic Coherence and Proper Scoring Rules." *IEEE Transactions on Information Theory* 55 (10): 4786–92.

Schervish, M. J., J. B. Kadane, and T. Seidenfeld. 2003. "Measures of Incoherence: How Not to Gamble If You Must." In *Bayesian Statistics 7: Proceedings of the 7th Valencia Conference on Bayesian Statistics*, ed. J. M. Bernardo, 385–402. Oxford: Clarendon.

Schervish, M. J., T. Seidenfeld, and J. B. Kadane. 2000. "How Sets of Coherent Probabilities May Serve as Models for Degrees of Incoherence." *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems* 8 (3): 347–55.

———. 2002. "Measuring Incoherence." *Sankhyā: The Indian Journal of Statistics* A 64:561–87.

Staffel, J. 2015a. "Disagreement and Epistemic Utility-Based Compromise." *Journal of Philosophical Logic* 44 (3): 273–86.

———. 2015b. "Measuring the Overall Incoherence of Credence Functions." *Synthese* 192 (5): 1467–93.

https://doi.org/10.1086/690715 Published online by Cambridge University Press

Wang, G., S. R. Kulkarni, H. V. Poor, and D. N. Osherson. 2011. "Aggregating Large Sets of Probabilistic Forecasts by Weighted Coherent Adjustment." *Decision Analysis* 8 (2): 128–44.
Williams, J. R. G. 2012. "Generalized Probabilism: Dutch Books and Accuracy Domination." *Journal of Philosophical Logic* 41 (5): 811–40.
Zynda, L. 1996. "Coherence as an Ideal of Rationality." *Synthese* 109 (2): 175–216.

https://doi.org/10.1086/690715 Published online by Cambridge University Press