Does non-measurability favour imprecision?

Cian Dorr

Draft of 3rd February 2023
(Final version to appear in *Mind*)

1 An argument for the permissibility of imprecise credence

Consider the following case, inspired by Isaacs, Hájek, and Hawthorne (2021: henceforth IHH):[1]

> *Strange Spinner:* God — who turns out to be quite the connoisseur of unusual games of chance — shows you an interesting contraption He has made. It consists of a pointer attached to the centre of a disc so that it is free to spin, with a lamp in the centre. He explains that He colored each point on the circumference either white or black, before covering the entire disc with a uniform grey coating. When the spinner is spun, the lamp will turn on if and only if the pointer ends up pointing to a white point. He chose the set of white points to be *non-measurable* with respect to the standard Lebesgue measure determined by angular distance on the disc's circumference.[2] In fact, it is as ill-behaved as possible as far as measurability is concerned: all its measurable subsets, and all its complement's measurable subsets, have Lebesgue measure 0.[3] He invites you to experiment with your spinner, adding that when you spin the pointer, it will randomly select a point on the circumference in a way that is *completely fair*. God then

[1] Hoëk (2021) has similar examples, though they are used for very different dialectical purposes.

[2] The *outer Lebesgue measure* $\lambda^*(X)$ of a set of n-tuples $X \subseteq [0,1)^n$ is defined as the greatest lower bound of $\{\sum_{n<\omega} m(B_n) \mid X \subseteq \bigcup_n B_n$ and each B_i is a box$\}$, where a *box* B is a product of n intervals and $m(B)$ is the product of those interval's lengths. Its *inner Lebesgue measure* $\lambda_*(X)$ can be defined as $1 - \lambda^*([0,1) - X)$. X is Lebesgue-measurable with Lebesgue measure $y = \lambda(X)$ when $y = \lambda^*(X) = \lambda_*(X)$. It is easy to show that the outer [inner] Lebesgue measure of a set is the greatest lower [least upper] bound of the Lebesgue measures of its measurable supersets [subsets]. We apply these functions to sets of points on the circumference of the spinner by arbitrarily choosing a point y and mapping each $x \in [0,1)$ to the point x rotations clockwise from y: since the functions are rotation-invariant, it makes no difference which y we choose.

[3] Equivalently, the set has inner Lebesgue measure 0 and outer Lebesgue measure 1.

> departs. You call after Him — 'Wait! Can You explain exactly what You mean by 'completely fair'?' — but He is gone. You wonder what doxastic attitude you should take to the proposition that the lamp will come on the first time you spin the spinner.

The question is hard because the case provides no principled considerations in favour of having credence x, for any real number $x \in [0, 1]$. Indeed, for each such x, the idea of responding to your situation by having credence x that the lamp will come on seems unsettlingly *arbitrary*, in a way that carries a whiff of irrationality.

Some philosophers theorize about doxastic states using a generalization of ordinary probability theory that seems tailor-made to help with puzzling cases like *Strange Spinner*. This theory posits two relations each of which can obtain between an agent, a proposition A, and at most one real number: the number can be the agent's *lower credence* in A, $\mathrm{Cr}_*(A)$, or the agent's *upper credence* in A, $\mathrm{Cr}^*(A)$. We always have $\mathrm{Cr}_*(A) \leq \mathrm{Cr}^*(A)$. When $\mathrm{Cr}_*(A) = \mathrm{Cr}^*(A)$, that number is one's *credence* in A, $\mathrm{Cr}(A)$. When $\mathrm{Cr}_*(A) < \mathrm{Cr}^*(A)$, one is said to be *credally imprecise* with respect to A; in that case, no real number is one's credence in A.[4]

IHH use the case in *Strange Spinner* to argue that credal imprecision is sometimes *rationally permissible*. Their argument implies that it is permissible for your lower and upper credences that the lamp will come on when you spin the spinner to be 0 and 1, respectively (assuming it is rationally permissible to have credence 1 in what God tells you).

Before getting into their argument, I should admit that I don't really understand its conclusion, since—in spite of the valiant efforts of a long philosophical tradition (Keynes 1921; Koopman 1940; Levi 1974; van Fraassen 1990; Walley 1991; Kaplan 1996; Joyce 2010; Sturgeon 2020)—I don't understand what it *is* to be credally imprecise with respect to a proposition. Various answers have been suggested, but they diverge so much that I am doubtful that there is any single phenomenon that theorists of credal imprecision have been talking about.

There are indeed some quite widely-accepted definitions of lower and upper credence, but they invoke other imprecision-theoretic terminology that raises the same interpretative challenges. A common approach (adopted by IHH) is to generalize the use of 'credence' so that we say that in some cases, one's credence in a proposition is not a real number, but rather a *set* of real numbers. One's lower and upper credences can then be defined as the greatest lower and least upper bounds of this set. If anything, this just exacerbates the difficulty. When $x < y$,

[4] In assuming credal imprecision to be incompatible with real-valued credence, I am ruling out a vagueness-theoretic interpretation which identifies being credally imprecise with respect to A with being such that no real number is *definitely* one's credence in A. Given classical logic, we cannot take 'no F is definitely G' to be incompatible with 'some F is G' without trivializing 'definitely'. Some fans of imprecision (van Fraassen 1990; Rinard 2013) favour a "supervaluationist interpretation" which suggests this identification; but IHH (§6) firmly reject it. I will explain in §3.2 how my argument can be adapted to fit the vagueness-theoretic interpretation.

many different sets have lower bound x and upper bound y; thus even if we had a good account of what it would be for one's lower and upper credences in A to be x and y, it is not obvious that it would extend to an account of what it would be for one rather than another of the sets with greatest lower bound x and least upper bound y to be one's credence in A.[5] Harder still are the questions raised by the popular analysis of set-valued credence according to which each agent has a *credal representor*, a set of probability functions, such that one's credence in A is the set containing exactly the numbers to which some member of one's credal representor maps A.[6] Even if we had a good account of what it is to have a certain set-valued credence function, it is not obvious that it would extend to an account of what it is to have a certain credal representor, since two distinct sets of probability functions can agree as regards the set to which they map every proposition.[7] For our present purposes, the minimally expressive concepts of lower and upper credence will suffice.

More informative analyses of imprecise credence are certainly possible, though much more controversial. One simple candidate is what I'll call the "inner-outer analysis", explored by Fagin and Halpern (1991). According to it, one's lower credence in A is one's *inner* credence in A: that is, the greatest lower bound of the set of real numbers that are one's credences in propositions that entail A. Likewise, one's upper credence in A is one's *outer* credence in A: the least upper bound of one's credences in propositions that A entails. The inner-outer analysis is appealingly simple, but has some consequences that few proponents of imprecise credence accept. For example, it implies that anyone who assigns a credence to one necessity and one impossibility automatically has lower and upper credences in every proposition. It also implies that if the range of one's credence function is a finite set, then all one's lower and upper credences must belong to it.[8] So I will not assume it, or any other analysis of lower and upper credence.

IHH in fact offer several lines of argument for their thesis that credal imprecision is rationally permissible in cases like *Strange Spinner*, but I will mostly be

[5] It might be proposed that only closed intervals $[x, y]$ can be credences: this makes the set-valued terminology a notational variant of the terminology of lower and upper credences.

[6] See Bradley 2019 and the many references therein. IHH express some scepticism about this framework.

[7] For example, suppose we have three mutually exclusive, jointly exhaustive, contingent propositions A, B, C. Let X be the set of all probability functions defined on truth-functional combinations of these propositions; let Y be the set of all such probability functions that assign equal probability to two of the three. Both X and Y assign $[0, 1]$ to A, B, C, $A \vee B$, $A \vee C$, and $B \vee C$, which are (up to necessary equivalence) the only contingent propositions in their domain.

[8] Fagin and Halpern (1991) show that the inner-outer analysis has another controversial consequence, namely that your unconditional lower credences obey the "complete monotonicity" conditions characteristic of Dempster-Shafer belief functions (Shafer 1976): a special case is that $\mathrm{Cr}_*(A \vee B) \geq \mathrm{Cr}_*(A) + \mathrm{Cr}_*(B) - \mathrm{Cr}_*(A \wedge B)$. For an argument that rational lower credences need not obey this constraint, see Walley 1991: §5.13.4.

focusing on one. (I'll discuss others in §5.) The argument I'm concerned with turns on another potentially controversial notion, that of *imprecise objective chance*. We assume that at a time t, a proposition A can have an *lower chance* $\mathrm{Ch}_{t*}(A)$ and an *upper chance* $\mathrm{Ch}_t{}^*(A)$, where $0 \leq \mathrm{Ch}_{t*}(A) \leq \mathrm{Ch}_t{}^*(A) \leq 1$. If the two are equal, this number is also $\mathrm{Ch}_t(A)$, A's *chance* at t; otherwise, no real number is A's chance at t.

Those who understand objective chance primarily in terms of its relationship with rational credence will plausibly be able to leverage their understanding of imprecise credence to make sense of imprecise chances; but such conceptions of objective chance are of course controversial. On the other hand, the inner-outer analysis of imprecise chance (on which a proposition's lower and upper chance are its inner and outer chances, i.e. the least upper bound of the chances of propositions that entail it and the greatest lower bound of the chances of propositions it entails) looks more defensible than the inner-outer analysis of imprecise credence. For example, it is less clear why one would think it possible for a proposition to lack a lower or upper chance, or for the truth about lower and upper chances to be interesting in a world where the truth about real-valued chance is uninteresting.[9]

To express the principle that drives IHH's argument, we will also need to avail ourselves of the notion of *conditional* imprecise credence. By contrast with the familiar ratio formula $\mathrm{Cr}(A \mid B) = \mathrm{Cr}(A \wedge B)/\mathrm{Cr}(B)$ for real-valued conditional credence, there is no natural candidate definition of conditional lower or upper credence in terms of the unconditional notions. Nevertheless, the conditional concepts seem no more obscure than the unconditional ones, and they are helpful for the purposes of systematic theorizing (cf. Hájek 2003).[10] Once we have the conditional notions, we can harmlessly define unconditional lower and upper credence in terms of them, as lower and upper credence conditional on some tautology $\top$.[11]

Once we have the concepts of imprecise chance and conditional imprecise

[9]IHH seem to endorse the inner-outer analysis of imprecise chance, although it is not needed for their argument.

[10]The inner-outer analysis in the form given above does not generalize straightforwardly to a definition of conditional lower and upper credence in terms of conditional real-valued credence. But Ruspini (1987) shows that, assuming one's credence function Cr is probabilistically coherent, the inner-outer analysis is equivalent to defining $\mathrm{Cr}_*(A)$ and $\mathrm{Cr}^*(A)$ as the greatest lower and least upper bounds of $\{\mathrm{Pr}(A) \mid \mathrm{Pr}$ is a probabilistically coherent extension of Cr defined on $A\}$. This reformulation does generalize straightforwardly: $\mathrm{Cr}_*(B \mid A)$ and $\mathrm{Cr}^*(B \mid A)$ can be defined as the limits of $\{\mathrm{Pr}(B \mid A) \mid \mathrm{Pr}$ is a probabilistically coherent extension of Cr, defined on A and B and such that $\mathrm{Pr}(A) > 0\}$.

[11]Some might think that an imperfectly rational agent could have an unconditional lower or upper unconditional credence in some proposition distinct from their lower or upper conditional credence in it given $\top$. But we can ignore such possibilities, since we are concerned with the credences one *should* have.

credence, we can use them to state a generalization of the "Principal Principle" (Lewis 1980). That principle says that one's credences, conditional on a hypothesis about the chances, should match the chances ascribed by that hypothesis; the generalization says something analogous about lower and upper credences and lower and upper chances. It can be stated as follows:

> Gen-PP If Cr_* and Cr^* are a rational combination of a lower and upper prior credence function and E is admissible at t, then for any proposition A:
>
> $$\mathrm{Cr}_*(A \mid E \wedge \mathrm{Ch}_{*t}(A) = x \wedge \mathrm{Ch}^*{}_t(A) = y) = x$$
> $$\mathrm{Cr}^*(A \mid E \wedge \mathrm{Ch}^*{}_t(A) = x \wedge \mathrm{Ch}^*{}_t(A) = y) = y$$
>
> so long as these are defined.

'E is admissible at t' can be glossed (at least as a first pass) as 'E is entirely about history up to and including t, including the chances at times up to and including t' (see Lewis 1980).[12]

To apply Gen-PP to *Strange Spinner*, we will need to cash out the "fairness" of the spinner in the terminology of lower and upper chance. This is naturally done as follows. Where X is any subset of $[0, 1)$ and y is any point on the circumference of the disc, say that X is *y-selected* iff the next time the spinner is spun, the pointer will point to a point whose clockwise angular distance from y (in revolutions) belongs to X. Then we take the fairness claim to imply that for any y, the lower and upper chances that X is y-selected are respectively $\lambda_*(X)$ and $\lambda^*(X)$, the Lebesgue inner and outer measures of X (see note 2). Note that this can consistently hold for all y, since the Lebesgue measure is translation-invariant. Since the set of white points has Lebesgue inner measure 0 and outer measure 1, and with chance 1 the lamp turns on iff the selected point is white, it follows that the lower and upper chances that the lamp will light when the spinner is first spun are 0 and 1, respectively.[13] Thus Gen-PP implies that conditional on the description of the case, you should have lower credence 0 and upper credence 1 that the lamp will light after the first spin.

IHH do not actually endorse Gen-PP. They do say that it 'is plausible in much the same way that the original [Principal] principle is', but are careful not to rule out a kind of extreme permissivism which would rule out both principles. Since they are only arguing for the *permissibility* of imprecise credence, the following weaker premise is all they need:

[12] IHH write their version of Gen-PP as $\mathrm{Cr}(A \mid \mathrm{Ch}(A) = [x, y]) = [x, y]$. This differs from mine in a few ways that are not relevant here: it concerns posterior credence, does not mention the admissible proposition E, and uses the ideology of set-valued credence and chance.

[13] I assume for simplicity that there is zero chance that it will never be spun.

Permissive Gen-PP It is always rationally permissible to have credences that conform to Gen-PP.

This entails the permissibility of credal imprecision, given the further premise that there are cases where one should have nonzero credence that a proposition has imprecise chance. I will grant—for the sake of argument, and *pace* Hoëk 2021—that there are such cases, and that *Strange Spinner* is one of them. But I will argue that Permissive Gen-PP is false.

2 Against Permissive Gen-PP

Let's continue the story:

> *Two Spinners:* The next day, God comes back with *two* identical-looking spinners. He explains that one—the *interesting* spinner—is the same one he showed you on before. The second—the *boring* spinner — is also designed such that the lamp will come on if the pointer ends up pointing to a white point. But on it, the white points form a contiguous interval comprising one-third of the disc's circumference. God then hands you one of the spinners, telling you that He made the choice by tossing a fair coin: if it came up Heads, you have the boring spinner, while if it came up Tails, you have the interesting spinner. As He departs, He reminds you that spinning the pointer will randomly select a point on the circumference in a way that is *completely fair*. Left alone with your spinner, you decide to experiment. The first time you spin it, the lamp turns on after the pointer comes to rest. The second time, it stays off. You keep at it, making records as you go: a 1 when the lamp comes on, a 0 when it doesn't. After thirty spins, your tally looks like this:
>
> 101011010001010000000100100010
>
> That's ten 1s and twenty 0s: just the proportion that would be most likely if you had the boring spinner. You wonder how you should react.

I say you should become *more confident* that you have the boring spinner than you were at the outset.

I hope readers will share my sense that this is obviously the right reaction. For those who feel the need of an argument, I offer the following: after those observations, you should be *pretty confident* that you have the boring spinner; you should not be pretty confident initially; but necessarily, if you are pretty confident afterwards and not before, you are more confident afterwards than before.[14] If you

[14] For a defence of the final premise, see Dorr, Nebel, and Zuehl 2021.

still don't feel the pull, bear with me; I will say more in defence of the judgment in §3.4. First, let's see how it can be used to argue agains Permissive Gen-PP.

We will of course need some auxiliary premises. The first is a bridge principle linking the ordinary expression 'more confident than' with the theoretical jargon of imprecise credence. No such principle is beyond dispute: part of what makes the jargon difficult is the lack of uncontroversial connections between it and ordinary words like 'confident'. But the following rather weak principle will suffice for our purposes:

> Confidence If x at t_1 is more confident in A than y at t_2 is in B, then x's lower credence at t_1 in A is not less than y's lower credence at t_2 in B.

Given Confidence, our starting judgment implies that your lower credence that you have the boring spinner should not be lower after your observations than it was initially.[15]

The second auxiliary premise we will need concerns how imprecise credences should change in response to new evidence. Using such a principle, we can transform our claim about how your credal state should evolve over time into one about your credal state at the initial time. Since we are already helping ourselves to the concepts of *conditional* lower and upper credence, there is a very natural candidate principle of this sort, namely the imprecise analogue of orthodox Bayesian conditionalization:

> Generalized Conditionalization Suppose you are rational and acquire total evidence E between t and t^+. Then for any propositions A and B, your lower and upper credences at t^+ in A given B are equal, respectively, to your lower and upper credences at t in A given BE (provided these are all well-defined).

Given our assumption that having credence x in A is equivalenty to having x as both one's lower and upper credence in A conditional on $\top$, Generalized Conditionalization implies the standard Bayesian conditionalization principle for real-valued credence. Familiar objections to that principle thus also apply to Generalized Conditionalization. However, these objections are orthogonal to our present concerns. For example, there are worries about cases where one's memory might fail (Arntzenius 2003), and worries related to the difficulty of finding an interpretation of "evidence" on which it's plausible both that one should have credence 1 in one's evidence and that one's credential state should

[15]Confidence can be strengthened without appreciable loss of plausibility to the claim that if x at t_1 is *at least* as confident in A as y is at t_2 in B, then x's lower credence at t_1 in A is greater than or equal to y's lower credence at t_2 in B *and x's upper credence at t_1 in A is greater than or equal to y's upper credence at t_2 in B*. The converse of this conditional also has some appeal, though it is far more controversial.

only change when one's evidence does (Jeffrey 1965). While these worries are important, they do not impugn the use of standard conditionalization (at least as an approximation) in examples where it is clear what proposition should play the role of one's "total evidence". In the imprecise setting, it is plausible that Generalized Conditionalization deserves a similar status.

Given Confidence and Generalized Conditionalization, our starting judgment implies that initially, your unconditional lower credence in the proposition that you have the boring spinner should not exceed your lower credence in that proposition conditional on the proposition that the first 30 spins have the given sequence of outcomes. To see why this is problematic for Permissive Gen-PP, we will need to further flesh out our interpretation of God's claim about the "fairness" of your spinner. In the previous section, we took that claim to imply that for any $X \subseteq [0,1)$ and any point y on the spinner's circumference, the inner and outer chances that the next spin will y-select X (i.e. point to a point whose clockwise angular distance from y belongs to X) equal the Lebesgue inner and outer measures of X. Now we must somehow extend this to propositions about where the pointer lands after the first n spins. Since the Lebesgue inner and outer measures λ_* and λ^* can be defined on all subsets of $[0,1)^n$, there is a very natural extension. Say that $X \subseteq [0,1)^n$ is y-selected at t iff X contains the n-tuple of numbers corresponding to the angular distances from y of the points where the pointer will land the next n times it is spun after t. Then the fairness claim is naturally understood to imply that the lower and upper chances at t of X being y-selected are respectively $\lambda_*(X)$ and $\lambda^*(X)$.

Importantly, this interpretation secures an imprecise analogue of probabilistic independence for distinct spins. Just as the Lebesgue measure of the Cartesian product of some sets is the product of their individual Lebesgue measures, so the Lebesgue inner and outer measures of the Cartesian product of some sets are products of their individual Lebesgue inner and outer measures: $\lambda_*(X_1 \times \cdots \times X_n) = \lambda_*(X_1) \cdots \lambda_*(X_n)$ and $\lambda^*(X_1 \times \cdots \times X_n) = \lambda^*(X_1) \cdots \lambda^*(X_n)$.[16] So the lower and upper

[16]We give the proof for inner measure; the case of outer measure is analogous. We need the following lemma: for every $X \subseteq [0,1)^k$, there is a Lebesgue-measurable $Y \subseteq X$ such that $\lambda_k(Y) = \lambda_{k*}(X)$. (Proof: By definition of λ_{k*}, for every $\varepsilon > 0$ there is a $Y_\varepsilon \subseteq X$ such that $\lambda_k(Y_\varepsilon) > \lambda_{k*}(X) - \varepsilon$; so we can take $Y = \bigcup_n Y_{1/n}$.) So, let $Y_1, \ldots, Y_n$ be Lebesgue-measurable subsets of $[0,1)$ such that $Y_i \subseteq X_i$ and $\lambda(Y_i) = \lambda_*(X_i)$, and let Z be a Lebesgue-measurable subset of $[0,1)^n$ such that $Z \subseteq \prod_i X_i$ and $\lambda_n(Z) = \lambda_{n*}(\prod_i X_i)$. For any $X \subseteq [0,1)$, let $C^n_i[X]$ denotes the set of n-tuples in $[0,1)^n$ whose ith co-ordinate is in X. Since $\prod_i X_i = \bigcap_i C^n_i[X_i]$, we have $Z \subseteq C^n_i[X_i]$ for each i. Also, since $Y_i \subseteq X_i$, $C^n_i[Y_i] \subseteq C^n_i[X_i]$. Thus $Z \cup C^n_i[Y_i] \subseteq C^n_i[X_i]$, and so $\lambda_n(Z \cup C^n_i[Y_i]) \leq \lambda_{n*}(C^n_i[X_i]) = \lambda_*(X_i) = \lambda(Y_i) = \lambda_n(C^n_i[Y_i])$. Since we also obviously have $\lambda_n(C^n_i[Y_i]) \leq \lambda_n(Z \cup C^n_i[Y_i])$, we can conclude that $\lambda_n(Z \cap \overline{C^n_i[Y_i]}) = 0$ for every i. Thus $\lambda_n(Z \cap \overline{\prod_i Y_i}) = \lambda_n(Z \cap \overline{\bigcap_i C^n_i[Y_i]}) = \lambda_n(Z \cap \bigcup_i \overline{C^n_i[Y_i]}) = \lambda_n(\bigcup_i (Z \cap \overline{C^n_i[Y_i]}) \leq \sum_i \lambda_n(Z \cap \overline{C^n_i[Y_i]}) = 0$, and hence $\lambda_{n*}(\prod_i X_i) = \lambda_n(Z) = \lambda_n(Z \cap \prod_i Y_i) \leq \lambda_n(\prod_i Y_i)$. But of course we also have $\lambda_n(\prod_i Y_i) \leq \lambda_{n*}(\prod_i X_i)$ (since each $Y_i \subseteq X_i$ and thus $\prod_i Y_i \subseteq \prod_i X_i$). We can conclude that $\lambda_{n*}(\prod_i X_i) = \lambda_n(\prod_i Y_i) = \prod_i \lambda(Y_i) = \prod_i \lambda_*(X_i)$, using the fact that the Lebesgue measure of a

chance at t of the proposition that $X_1 \times \cdots \times X_n$ is y-selected at t—i.e., that the first spin after t will y-select X_1, and ..., and the nth spin will y-select X_n—are respectively $\lambda_*(X_1)\cdots\lambda_*(X_n)$ and $\lambda^*(X_1)\cdots\lambda^*(X_n)$, the products of the lower and upper chances of the propositions that the ith spin will y-select X_i.

This interpretation provides chance-theoretic claims that can be plugged into Gen-PP to derive claims about the lower and upper conditional credences you should have in various outcomes, conditional on having the boring and interesting spinners; Permissive Gen-PP will yield corresponding permissibility claims. However, these claims do not immediately yield what we are looking for, namely a value for your lower conditional credence that you have the boring spinner conditional on the given sequence of outcomes. To derive such a claim, we will need to appeal to some basic *coherence principles* about lower and upper conditional credence.

We will need two families of such principles, each generalizing a familiar coherence principle for real-valued conditional probablility . The first family generalizes the following *multiplicative axiom*, basic to the theory of real-valued conditional probability:

$$\text{(M)} \qquad \mathrm{Cr}(AB \mid C) = \mathrm{Cr}(A \mid BC)\,\mathrm{Cr}(B \mid C)$$

When C is a tautology and $\mathrm{Cr}(B) > 0$, this is equivalent to the familiar ratio formula $\mathrm{Cr}(A \mid B) = \mathrm{Cr}(AB)/\mathrm{Cr}(B)$. In theories like that of K. Popper (1959) and K. R. Popper (1955) that treat conditional probability as primitive, (M) plays a fundamental role in co-ordinating probabilities conditional on different propositions.[17] Proponents of imprecise credence need some principles that can do a similar job. And in fact, all extant formal treatments of imprecise conditional probability imply the following generalizations of (M):

$$\begin{array}{llll}
(\mathrm{M}_{***}) & \mathrm{Cr}_*(AB|C) \geq \mathrm{Cr}_*(A|BC)\,\mathrm{Cr}_*(B|C) & (\mathrm{M}^{***}) & \mathrm{Cr}^*(AB|C) \leq \mathrm{Cr}^*(A|BC)\,\mathrm{Cr}^*(B|C) \\
(\mathrm{M}_{**}{}^{*}) & \mathrm{Cr}_*(AB|C) \leq \mathrm{Cr}_*(A|BC)\,\mathrm{Cr}^*(B|C) & (\mathrm{M}^{**}{}_{*}) & \mathrm{Cr}^*(AB|C) \geq \mathrm{Cr}^*(A|BC)\,\mathrm{Cr}_*(B|C) \\
(\mathrm{M}_{*}{}^{*}{}_{*}) & \mathrm{Cr}_*(AB|C) \leq \mathrm{Cr}^*(A|BC)\,\mathrm{Cr}_*(B|C) & (\mathrm{M}^{*}{}_{*}{}^{*}) & \mathrm{Cr}^*(AB|C) \geq \mathrm{Cr}_*(A|BC)\,\mathrm{Cr}^*(B|C)
\end{array}$$

For example, all six inequalities fall out of the "representor" approach. Since each member of the representor is (or determines via the ratio formula) a conditional probability function obeying (M), there is no way that the greatest lower bound of the members' probabilities in AB given C could, for example, be lower than the product of the greatest lower bounds of their credences in A given BC

product is the product of Lebesgue measures.

[17](M) also holds in Kolmogorov's theory of "regular conditional probability" (see Easwaran 2019), where claims of conditional probability must always be relativized to a *sub-σ-algebra*, provided that we interpret all three conditional probabilities as relativized to the same sub-σ-algebra.

and in B given C. Other analyses of lower and upper conditional probability (e.g. that of Walley 1991: based on a concept of 'prevision' for random variables) also concur in validating these inequalities. Together, they do quite a lot to constrain $\mathrm{Cr}_*(\cdot \mid BC)$ and $\mathrm{Cr}^*(\cdot \mid BC)$ given $\mathrm{Cr}_*(\cdot \mid C)$ and $\mathrm{Cr}^*(\cdot \mid C)$, though (by contrast with the real-valued case) they do not completely pin them down, even when $\mathrm{Cr}_*(B \mid C) > 0$.

In the case one of the three relevant pairs of propositions has a real-valued conditional probability (i.e. has equal lower and upper conditional probabilities), we can combine pairs of inequalities to derive some useful equations:

$$
\begin{array}{llll}
(\mathrm{M}_{**}\cdot) & \mathrm{Cr}_*(AB|C) = \mathrm{Cr}_*(A|BC)\,\mathrm{Cr}(B|C) & (\mathrm{M}^{**}\cdot) & \mathrm{Cr}^*(AB|C) = \mathrm{Cr}^*(A|BC)\,\mathrm{Cr}(B|C) \\
(\mathrm{M}_*\cdot_*) & \mathrm{Cr}_*(AB|C) = \mathrm{Cr}(A|BC)\,\mathrm{Cr}_*(B|C) & (\mathrm{M}^*\cdot^*) & \mathrm{Cr}^*(AB|C) = \mathrm{Cr}(A|BC)\,\mathrm{Cr}^*(B|C) \\
(\mathrm{M}\cdot^*{}_*) & \mathrm{Cr}(AB|C) = \mathrm{Cr}^*(A|BC)\,\mathrm{Cr}_*(B|C) & (\mathrm{M}\cdot_*{}^*) & \mathrm{Cr}(AB|C) = \mathrm{Cr}_*(A|BC)\,\mathrm{Cr}^*(B|C)
\end{array}
$$

In each case, the two inequalities got by replacing the bullet with an upper or lower asterisk imply that the equality holds whenever Cr is defined on the relevant two propositions (i.e., whenever Cr_* and Cr^* agree on them).[18]

The second family of coherence principles we will need generalize the following *finite additivity* axiom from the theory of real-valued conditional probability:

$$
(\mathrm{A}) \qquad \mathrm{Cr}(A \mid C) = \mathrm{Cr}(AB \mid C) + \mathrm{Cr}(A\overline{B} \mid C)
$$

This plays a similarly fundamental role to (M).[19] The following inequalities can play an analogous role in the theory of imprecise conditional credence:[20]

$$
\begin{array}{llll}
(\mathrm{A}_{***}) & \mathrm{Cr}_*(A|C) \geq \mathrm{Cr}_*(AB|C) + \mathrm{Cr}_*(A\overline{B}|C) & (\mathrm{A}^{***}) & \mathrm{Cr}^*(A|C) \leq \mathrm{Cr}^*(AB|C) + \mathrm{Cr}^*(A\overline{B}|C) \\
(\mathrm{A}_{**}{}^*) & \mathrm{Cr}_*(A|C) \leq \mathrm{Cr}_*(AB|C) + \mathrm{Cr}^*(A\overline{B}|C) & (\mathrm{A}^{**}{}_*) & \mathrm{Cr}^*(A|C) \geq \mathrm{Cr}^*(AB|C) + \mathrm{Cr}_*(A\overline{B}|C) \\
(\mathrm{A}_*{}^*{}_*) & \mathrm{Cr}_*(A|C) \leq \mathrm{Cr}^*(AB|C) + \mathrm{Cr}_*(A\overline{B}|C) & (\mathrm{A}^*{}_*{}^*) & \mathrm{Cr}^*(A|C) \geq \mathrm{Cr}_*(AB|C) + \mathrm{Cr}^*(A\overline{B}|C)
\end{array}
$$

These follow from the "representor" analysis for exactly same reason as the analogous multiplicative inequalities; they also fall out from other approaches such as that of Walley (1991). And as before, pairs of inequalities can be combined to yield useful equations in cases where one of the three proposition-pairs receives

[18]These equations entail in turn that (M) holds whenever all three of $\mathrm{Cr}(B \mid C)$, $\mathrm{Cr}(AB \mid C)$, and $\mathrm{Cr}(A \mid BC)$ exist. Moreover, if $\mathrm{Cr}(A \mid BC)$ and one of $\mathrm{Cr}(B \mid C)$ and $\mathrm{Cr}(AB \mid C)$ exist, the other must also exist; and if both $\mathrm{Cr}(B \mid C)$ and $\mathrm{Cr}(AB \mid C)$ exist and $\mathrm{Cr}(B \mid C) > 0$, $\mathrm{Cr}(A \mid BC)$ must exist.

[19]K. Popper (1959) makes an exception to (A) to allow for special C with the property that $\mathrm{Cr}(A \mid C) = 1$ for all A; I prefer to let $\mathrm{Cr}(\cdot \mid C)$ go undefined in such cases.

[20]This presentation is redundant: the two inequalities on the last line follow from the ones above them by substituting $\overline{B}$ for B and appealing to double negation elimination.

a real-valued conditional credence:

$$(\mathrm{A}_{**}\cdot)\quad \mathrm{Cr}_*(A|C) = \mathrm{Cr}_*(AB|C) + \mathrm{Cr}(A\overline{B}|C) \qquad (\mathrm{A}^{**}\cdot)\quad \mathrm{Cr}^*(A|C) = \mathrm{Cr}^*(AB|C) + \mathrm{Cr}(A\overline{B}|C)$$

$$(\mathrm{A}_*\cdot_*)\quad \mathrm{Cr}_*(A|C) = \mathrm{Cr}(AB|C) + \mathrm{Cr}_*(A\overline{B}|C) \qquad (\mathrm{A}^*\cdot^*)\quad \mathrm{Cr}^*(A|C) = \mathrm{Cr}(AB|C) + \mathrm{Cr}^*(A\overline{B}|C)$$

$$(\mathrm{A}\cdot^*{}_*)\quad \mathrm{Cr}(A|C) = \mathrm{Cr}^*(AB|C) + \mathrm{Cr}_*(A\overline{B}|C) \qquad (\mathrm{A}\cdot_*{}^*)\quad \mathrm{Cr}(A|C) = \mathrm{Cr}_*(AB|C) + \mathrm{Cr}^*(A\overline{B}|C)$$

These give us back (A) in the case where all three (or indeed any two) of $\mathrm{Cr}(A \mid C)$, $\mathrm{Cr}(AB \mid C)$ and $\mathrm{Cr}(A\overline{B} \mid C)$ exist.

We now have all we need to complete our argument against Permissive Gen-PP. Assume you start out with the imprecise credences mandated by Gen-PP. Let B be *you got the boring spinner*; for the sake of generality, let E be some proposition that specifies the outcome of the first n spins for some $n \geq 1$, and entails that the lamp stayed off m times. (In the case we have been focusing on, $n = 30$ and $m = 20$.) Applying Gen-PP to the time before the coin-toss, we have

$$\mathrm{Cr}(B) = 1/2$$

Applying it to the time after the coin toss conditional on your getting the boring spinner, we have

$$\mathrm{Cr}(E \mid B) = 2^m/3^n$$

(Note that for large n, this is a very small no matter what m is.) More interestingly, we can apply Gen-PP to the case where you get the interesting spinner to derive the equations

$$\mathrm{Ch}_*(E \mid \overline{B}) = 0$$
$$\mathrm{Ch}^*(E \mid \overline{B}) = 1$$

For since the subset W of $[0,1)$ corresponding to the white points has inner Lebesgue measure 0 and outer Lebesgue measure 1, its complement does too. Hence any subset $X_1 \times \cdots \times X_n$ of $[0,1)^n$ where each X_i is either W or its complement has inner measure 0 (the product of the inner measures of $X_1, \ldots, X_n$) and outer measure 1 (the product of the outer measures of $X_1, \ldots, X_n$).

Given (M), $(\mathrm{M}_{**}\cdot)$ and $(\mathrm{M}^{**}\cdot)$, the above equations imply:

$$\begin{aligned}
\mathrm{Cr}(BE) &= \mathrm{Cr}(E \mid B)\,\mathrm{Cr}(B) = 2^m/3^n \times 1/2 = 2^{m-1}/3^n \\
\mathrm{Cr}_*(\overline{B}E) &= \mathrm{Cr}_*(E \mid \overline{B})\,\mathrm{Cr}(\overline{B}) = \quad 0 \times 1/2 = 0 \\
\mathrm{Cr}^*(\overline{B}E) &= \mathrm{Cr}^*(E \mid \overline{B})\,\mathrm{Cr}(\overline{B}) = \quad 1 \times 1/2 = 1/2
\end{aligned}$$

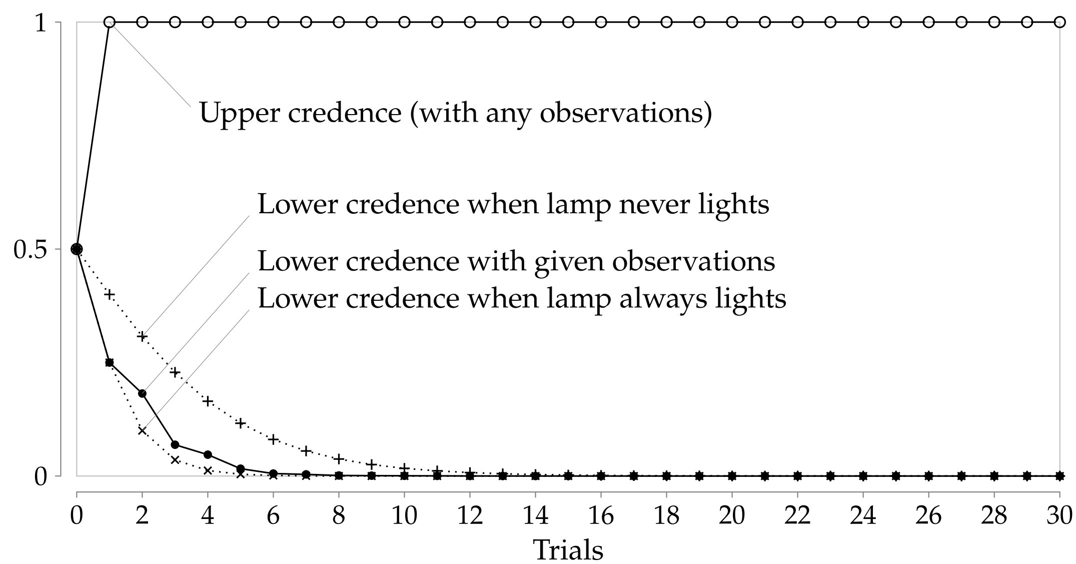

Figure 1: Evolution of your attitude to the proposition that you have the boring spinner

By ($\mathrm{A}_*\cdot_*$) and ($\mathrm{A}^*\cdot^*$) (setting $A := E$, $B := \overline{B}$, $C := \top$), we get

$$\begin{aligned} \mathrm{Cr}_*(E) &= \mathrm{Cr}(BE) + \mathrm{Cr}_*(\overline{B}E) = 2^{m-1}/3^n + 0 \quad = 2^{m-1}/3^n \\ \mathrm{Cr}^*(E) &= \mathrm{Cr}(BE) + \mathrm{Cr}^*(\overline{B}E) = 2^{m-1}/3^n + 1/2 = (2^m + 3^n)/(2 \times 3^n) \end{aligned}$$

These give us what we need to determine your conditional upper and lower credences in B given E by applying ($\mathrm{M}\cdot^*{}_*$) and ($\mathrm{M}\cdot_*{}^*$) (setting $A := B$, $B := E$, $C := \top$):

$$\begin{aligned} \mathrm{Cr}_*(B \mid E) &= \frac{\mathrm{Cr}(BE)}{\mathrm{Cr}^*(E)} = \frac{2^{m-1}/3^n}{(2^m + 3^n)/(2 \times 3^n)} = \frac{2^m}{2^m + 3^n} \\ \mathrm{Cr}^*(B \mid E) &= \frac{\mathrm{Cr}(BE)}{\mathrm{Cr}_*(E)} = \frac{2^{m-1}/3^n}{2^{m-1}/3^n} = 1 \end{aligned}$$

Given Generalized Conditionalization, it follows that your upper credence that you have the boring spinner will jump from 1/2 to 1 after your first observation and stay at 1 thereafter. Meanwhile, your lower credence that you have the boring spinner will decrease monotonically from its starting value of 1/2. Figure 1 depicts what this looks like for the sequence of outcomes given in the vignette, as well as a sequence in which the lamp comes on every time and one where it never comes on. Whatever you see, your lower credence will approach 0 as you make more observations. It will never be 1/2 or above. Given Confidence, this means that no observations you could make would make you more confident that you have the boring spinner.

I conclude that you should *not*, in this case, have the credences mandated by Gen-PP: thus Permissive Gen-PP is false.[21]

3 Responses

This section will discuss four possible responses to the previous section's argument. The first is to reject my interpretation of the case; the second is to reject Confidence; the third is to reject Generalized Conditionalization; and the fourth is to deny the starting judgment that your observations should make you more confident that you have the boring spinner.

3.1 *Rejecting the interpretation of the case*

In conversation, the most popular response has been to reject my interpretation of the "fairness" of the spinners, on which for any $X \subseteq [0,1)^n$ and any reference point y, the lower and upper chances of X being y-selected equal the Lebesgue inner and outer measures of X. It has been suggested that the the lower and upper chances of a sequence of lamp on/lamp off outcomes on the interesting spinner should be instead be determined from a representor which contains all and only the probability functions that treat the spins of the interesting spinner as probabilistically independent events with some fixed chance. To see the difference between the two interpretations, suppose you get the interesting spinner and consider the proposition that the lamp comes on after the first spin and not after the second. On my interpretation, its inner and outer chances are 0 and 1. On the alternative interpretation, its inner chance is still 0, but its outer chance is only 1/4, since the member of the representor that maximizes its probability is the one that treats the spins as independent events with probability 1/2.

So far, this proposal is rather ad hoc, since it only concerns the chances of propositions about when the lamp does and doesn't come on. But there is a natural way of generalizing it to arbitrary propositions about where the spinner lands after each of the first n spins. We can determine lower and upper chances for any such proposition using a representor that contains all and only the total (finitely additive) probability functions Pr which are such that, for some point y and

[21] Evolutions where one's lower credence decreases while one's upper credence increases are called *dilations*. White (2010) (building on Seidenfeld and Wasserman 1993) considers a case where proponents of imprecise probability counsel dilation with respect to the proposition that a coin landed heads, and deems this implausible. IHH's reply to White strikes me as convincing. But my reasons for thinking you should not dilate in *Two Spinners* are unrelated to White's reasons for opposing dilation in his case: his worry has to do with the *predictability* of dilation, whereas I simply think that your lower credence should increase rather than decrease.

some total (finitely additive) extension λ^+ of Lebesgue measure to the powerset of $[0, 1)$, $\Pr(X_1 \times \cdots \times X_n \text{ is } y\text{-selected}) = \lambda^+(X_1) \cdots \lambda^+(X_n)$, for all $X_1, \ldots, X_n \subseteq [0, 1)$. Alternatively, we could limit the representor to probability functions associated in this way with extensions of Lebesgue-measure that are both total and *rotation-invariant*.[22]

For several reasons, this does not seem to me to be a promising line of response. First, a key desideratum for IHH is that the lower and upper chances (and the corresponding credences) should respect relevant symmetries of the setup. In our case, such symmetries include not just *global* rotations which apply to every spin, corresponding to functions on $[0, 1)^n$ that map $\langle x_1, \ldots, x_n \rangle$ to $\langle x_1 + y, \ldots, x_n + y \rangle$ for some y, (where all addition is mod 1), but also *local* rotations which apply to only one spin, mapping $\langle x_1, \ldots, x_{i-1}, x_i, x_{i+1}, \ldots, x_n \rangle$ to $\langle x_1, \ldots, x_{i-1}, x_i + y, x_{i+1}, \ldots, x_n \rangle$. Since the spinner doesn't have a memory, insofar as it doesn't matter which point we assign to zero when mapping from $[0, 1)$ to propositions about a single spin, it seems that it should also not matter if we choose a different zero point each time when mapping from $[0, 1)^n$ to propositions about the first n spins. But while the alternative lower and upper chances are invariant under global rotations, they are not invariant under local rotations.[23]

Second, insofar as the concepts of lower and upper chance are in good standing, we can elaborate the case to have God explicitly tell you that the lower and upper chances work in accordance with my interpretation. Or if we want something a bit less theory-laden, we could specify that if the coin lands heads, God will supply you with a whole series of interesting spinners, each of which you get to spin just once. No two of these spinners are duplicates, although the set of white points on each of them has inner Lebesgue measure 0 and outer Lebesgue measure 1. In this case, there are no grounds whatsoever for limiting the chance-representor in the relevant way—for all you know, the pattern on the second spinner is derived from the pattern on the first by swapping black and white. But these variations do little to diminish the intuitive plausibility of the claim that the given sequence of observations should make you more confident that you have the boring spinner.

Third, the argument against Permissive Gen-PP can anyway be made using

[22]This interpretation requires giving up the identification of lower and upper chance with inner and outer chance. For it implies that for $X \subseteq [0, 1)^n$, X has a real-valued chance of being y-selected only if X is Lebesgue-measurable, so that inner and outer chances still correspond to inner and outer Lebesgue measures. For example, the *outer* chance of *lamp on, then lamp off* is 1, although its *upper* chance is $1/4$.

[23]Let V be a Vitali set in $[0, 1)$, with outer measure 1, and let W be its image under rotation through some nonzero rational angle. Let A be the proposition that the next two spins will respectively y-select V and $\overline{W}$ (the complement of W). Since V is a Vitali set, $V \cap W = \emptyset$, i.e. $V \subseteq \overline{W}$. The upper chance that $V \times \overline{V}$ is y-selected is $1/4$, since this proposition's probability is maximized by a λ^+ for which $\lambda^+(V) = \lambda^+(\overline{V}) = 0.5$. But the upper chance that $V \times \overline{W}$ is y-selected is 1. For since $\lambda^*(V) = 1$, there is a total extension λ^+ of Lebesgue measure for which $\lambda^+(V) = 1$ and hence also $\lambda^+(\overline{W}) = 1$.

the alternative interpretation. Consider your lower and upper credences after you have made 29 spins and seen the lamp come on ten times. The member of the proposed chance-representor that maximizes the probability of *that* outcome is the one that treats spins of the interesting spinner as independent events with probability 10/29. This probability function treats the observed outcome as slightly favouring the hypothesis that you have the interesting spinner. Thus after 29 spins, your lower credence that you have the boring spinner will be slamply below 1/2. Given Confidence, this rules out your being more confident that you have the boring spinner.

3.2 *Rejecting Confidence*

Another possible response involves accepting the judgment that you should get more confident that you have the boring spinner, while also accepting that your lower and upper credences should evolve in the way depicted in Figure 1. This requires rejecting Confidence.

Some discussions of the relation between imprecise credence and comparisons of confidence suggest independent motivations for giving up Confidence. Certain authors (van Fraassen 1990; Joyce 2010; Rinard 2013) posit a close connection between imprecise credence and *vagueness* in words like 'confident'. According to Rinard, for example, 'An agent is determinately more confident of P than Q if and only if every function $[Pr]$ in her representor has $Pr(P) > Pr(Q)$' (Rinard 2013). Since 'more confident' is vague, proponents of this biconditional should not also accept the corresponding biconditional without the 'determinately'. And insofar as one rejects this biconditional, Confidence looks unmotivated. If one does not rule out an agent's being more confident in A than B despite having a representor that contains probability functions for which $\Pr(A) \leq \Pr(B)$, why would one rule it out in the case where the agent has a higher lower credence in B than in A?

Rinard's biconditional does not fit naturally with the view we have been assuming (following IHH), on which being credally imprecise with respect to a proposition is inconsistent with there being a real number that is one's credence in it. It seems more suited to a view on which credal imprecision is only inconsistent with there being a unique real number that *determinately* one's credence in the relevant proposition; indeed, it even suggests that such imprecision might in fact entail that there *is* a unique real number that is one's credence in that proposition.[24] That alternative view weakens the prima facie case for Gen-PP,

[24]The latter claim would hold if the word 'credence' is vague in the following way: for any partial function Pr from propositions to reals, Pr is in one's representor iff it is not determinately false that (for all A and x, one has credence x in A iff $\Pr(A) = x$). One could also derive the claim from the view that for all Pr, Pr is in one's representor iff it is not determinately false that (for all A and B, $\Pr(A) \leq \Pr(B)$ iff one is at least as confident in B as in A): see the principle 'Credence Existence' in Dorr, Nebel, and Zuehl 2021.

since it suggests that the original Principal Principle does not in fact fall silent about the relevant cases. But it doesn't really matter. For the vagueness-theoretic objection to Confidence does not apply if we weaken that principle by adding 'determinately' to its antecedent, yielding:

> Weak Confidence If x is *determinately* more confident at t_1 that A than y is at t_2 that B, then x's upper credence at t_1 in A is greater than or equal to y's upper credence at t_2 in B, and x's lower credence at t_1 in A is greater than or equal to y's lower credence at t_2 in B.

This is enough for a variant of our argument against Permissive Gen-PP, since your observations in *Two Spinners* should plausibly leave you not only more confident, but *determinately* more confident, that you have the boring spinner. A response on which you should be more confident but not determinately more confident seems quite bizarre and unpalatable, given any reasonable foundational account of the phenomenon of vagueness.

Rejecting Weak Confidence would require a radical rethinking of the aspirations of the theory of imprecise credence. Bayesian epistemology has traditionally been centrally concerned with questions about confidence, e.g. whether certain observations should make you more or less confident in certain hypotheses. If even Weak Confidence is rejected, there seems to be nothing left to connect claims about how observations should affect your imprecise credal state with these classic questions. If we care about the confidence-theoretic questions, we will need some separate theory to answer them; it is unclear why we would even be interested in the theory of credal imprecision.

The strategy of rejecting Confidence (especially if it also involves rejecting Weak Confidence) tends moreover to exacerbate the challenge of explaining what it *is* to be credally imprecise, since it undercuts various explanations that turn on the familiar notion of confidence. Of course, there are approaches to this challenge that do not appeal to that notion, some of which seem prima facie friendly to failures of Confidence. For example, the time-honoured idea that credence can be understood in terms of *betting dispositions* has been developed by Walley (1991) into a versatile account of imprecise credence, where differences in lower and upper credence boil down to differences between the most one would pay to buy a certain gamble and the least one would accept to sell it.[25] Another time-honoured idea identifies credences with certain all-or-nothing *beliefs* or *judgments* about probability (in some broadly epistemic sense of 'probability'); this can also be extended to an account of credal imprecision either in terms of less-than-maximally-specific beliefs about probability, or positive beliefs about

[25]Mahtani (2016) suggests a different account of credal imprecision in terms of betting behaviour, on which it is characterized by a kind of *instability* in one's betting dispositions.

some independently understood notion of imprecise probability, or both.[26] Both of these ideas look well-suited to making sense of failures of CONFIDENCE, and maybe even WEAK CONFIDENCE. There is no *obvious* reason why some evidence could not render one (determinately) more confident in A while simultaneously decreasing the maximum one would pay for a one-dollar bet on A, or making one less opinionated about its epistemic probability, or even causing one to believe that its lower epistemic probability has decreased. But neither idea seems to do much to make the evolution prescribed by GEN-PP in *Two Spinners* more palatable. For after your observations, it seems intuitively reasonable for you to be disposed to pay considerably more than 50 cents for a bet that pays $1 if you have the boring spinner, and for you to sincerely make speeches like 'Probably, this is the boring spinner' or 'I think this is likely to be the boring spinner' (which plausibly express beliefs that the all-or-nothing theorist would take to be inconsistent with having lower credence less than 1/2). The challenge to WEAK GEN-PP can thus be raised not just from pre-theoretic judgments about *confidence*, but from pre-theoretic judgments about several other topics which one might wish to draw on in an explanation of credal imprecision.

3.3 *Rejecting GENERALIZED CONDITIONALIZATION*

A third possible response to the argument is to give up on GENERALIZED CONDITIONALIZATION. There is some precedent for this in the literature on imprecise credence. Several authors (Wilson 2001; Cattaneo 2008; Bradley 2022) have proposed that when you gain new evidence E, your new representor should be derived from your old one by first throwing away those probability functions that gave E especially low probabilities, and then conditionalizing each remaining probability function on E. This can easily lead to failures of GENERALIZED CONDITIONALIZATION.[27]

---

[26]IHH (§6) suggest that having lower credence x and upper credence y in a proposition might 'amount to the belief that it has inner probability x and outer probability y'. One might alternatively identify having lower and upper credences x y in A with being such that x and y are the infimum and supremum of $\{x' \mid \text{one believes that } A\text{'s lower epistemic probability is at least } x'\}$. This will make it possible to have lower and upper credences in A without being fully opinionated about A's lower and upper epistemic probabilities, and even while believing them to be identical.

[27]For example, suppose that you are to successively draw two balls from an urn. Each ball can be red or black, so there are four possibilities: R_1R_2, R_1B_2, B_1R_2, and B_1B_2. Your initial representor contains three probability functions:

	R_1R_2	R_1B_2	B_1R_2	B_1B_2
Pr_1	0.25	0.25	0.25	0.25
Pr_2	0.01	0	0	0.99
Pr_3	0	0.01	0.99	0

Thus your lower credence in R_2 conditional on R_1 is 0 (thanks to Pr_3) and your upper credence is 1, (thanks to Pr_2). But assuming the 0.01 value Pr_2 and Pr_3 assign to R_1 counts as "extreme" enough (relative to the 0.5 assigned by Pr_1) for them to be discarded after you observe R_1, this

However, we can specify favourable circumstances where the "throwing away" procedure these authors favour will agree with Generalized Conditionalization as regards the lower or upper credence of some proposition. Suppose you have a real-valued credence x in HE: i.e., $\Pr(HE) = x$ for every Pr in your representor. Then the members Pr of your representor that maximize $\Pr(E)$ are exactly those that maximize $\Pr(\overline{H}E)$, and hence also maximize $\Pr(\overline{H} \mid E)$, i.e. $\Pr(\overline{H}E)/(\Pr(\overline{H}E) + x)$.[28] Since these Pr will not be discarded after you learn E, we have the following:

> Weak Generalized Conditionalization Suppose you are rational and acquire total evidence E between t and t^+, and have a real-valued credence in HE at t. Then your upper credence at t^+ in $\overline{H}$ equals your upper credence at t in $\overline{H}$ conditional on E, and your lower credence at t^+ in H equals your lower credence at t in H conditional on E.

This principle is strong enough to substitute for Generalized Conditionalization in our argument. Taking H to be the proposition that you have the boring spinner and E to be some sequence of observed outcomes, the hypothesis of Weak Generalized Conditionalization is satisfied; so we can conclude that your lower credence that you have the boring spinner diminishes as depicted in Figure 1. This is enough, given Confidence, to rule out your becoming more confident. The "throwing away" procedure just blocks the argument that your *upper* credence goes to 1. Indeed, on natural precisifications of the view, sufficiently many observations will leave your upper credence *also* far below 1/2, as your representor comes to be dominated by the descendents of probability functions that initially assigned high probability to the interesting spinner generating the outcomes you observed. This plausibly implies that you will end up *less* confident that you have the boring spinner. Thus, far from helping with our example, the strategy of 'learning by ignoring the most wrong' (Bradley 2022) actually makes the problem worse.

Of course, this strategy is not the only possible way to give up Generalized Conditionalization, and some might hope to handle the case by developing some novel alternative updating method.[29] But the difficulties facing this project

observation will leave you with the singleton representor $\{\Pr_1(\cdot \mid R_1)\}$, and thus with credence 0.5 in R_2.

[28]Since for positive x, y, z, $\frac{y}{x+y} \leq \frac{z}{x+z}$ iff $xy + yz \leq xz + yz$, iff $xy \leq xz$, iff $y \leq z$, iff $x + y \leq x + z$.

[29]Weatherson (2007) defends a "dynamic Keynesianism" on which, as on the "throwing away" view discussed above, one should sometimes discarding some functions in the representor before conditionalizing the remainder on the evidence. On Weatherson's view, however, the functions to be discarded are identified not by the probability they assign the evidence, but by some mysterious factors for which he provides no formal model. This is supposed to model a special kind of "learning about fundamental evidential relationships". But it is not promising to attribute the failure of Generalized Conditionalization in *Two Spinners* to this supposedly special kind of learning. We can imagine our agent as already equipped with a detailed take on evidential relationships, formed during a long career of hypothesis-testing.

should not be underestimated. True, the most widely-discussed arguments for conditionalization—Dutch book arguments (Lewis 1999) and arguments relating to the value of accuracy (Greaves and Wallace 2006)—are tricky to generalize to imprecise credence, due to a lack of consensus about the decision-theoretic upshot of imprecise credences and about how to measure their accuracy.[30] But quite apart from these arguments, there is also a simple normative intuition that supports both ordinary conditionalization and its imprecise generalization: namely, that when you are trying to gather evidence to compare two hypotheses H_1 and H_2, you need evidence that at least one of the two fails to entail. If what you learn is entailed by both hypotheses, it is irrelevant to the comparison. This strikes me as a deeply plausible thought even if it is not backed up by any further argument. And it implies Generalized Conditionalization, assuming 'You learnt nothing relevant to the comparison between H_1 and H_2' implies 'Your lower and upper credences in H_1 and H_2 conditional on their disjunction should not change'.[31]

3.4 *Biting the bullet*

A final response is to reject the judgment that you should respond to your observations by becoming more confident that you got the boring spinner. This might be tolerable if the surprise could be limited to far-fetched thought experiments. Unfortunately, it generalizes in a disastrous way to many real-world inquiries.

To see the issue, note first that our analysis of *Two Spinners* wouldn't differ much if the choice of spinner was based on a heavily biased coin rather than a fair one. Whatever credences you start out with, your observations will cause your lower credence that you have the boring spinner to approach 0 just as fast as your *credence* would if you had initially been certain that the interesting spinner would produce exactly the outcomes you observed. This is very fast! For example, even if you initially had credence 0.9999999 that you had the boring spinner, any series of 30 observations will leave your lower credence less than 1/3. But this version of the case is quite similar to what we face whenever we try to investigate the laws governing some perhaps-indeterministic kind of physical system, such as the decay of free neutrons. If the concepts of lower and upper chance are in good standing, it seems unreasonably dogmatic to begin our inquiry with credence *zero* in hypotheses according to which the laws governing the relevant physical systems involve imprecise chances analogous to those generated by the interesting

[30] Although Walley (1991: ch. 6), defining imprecise credence in terms of betting dispositions, is able to derive Generalized Conditionalization using a kind of Dutch Book argument.

[31] By setting $H_1 = AE$ and $H_2 = E$, we can infer that your lower and upper conditional credences in A given E should not change when you learn E. But since we assume that learning E involves giving real-valued credence 1 to E, we know (by ($\mathrm{M}_{**}\cdot$), ($\mathrm{M}^{**}\cdot$), ($\mathrm{A}_{**}\cdot$), and ($\mathrm{A}^{**}\cdot$)) that after learning E, your lower and upper unconditional credences in A equal your lower and upper credences in A given E.

spinner.[32] But so long as we are open-minded enough to assign a positive upper prior credence to such a hypothesis, Gen-PP and Generalized Conditionalization imply that any observations we might make should cause our upper credences in it to rapidly approach 1, so that our lower credences in all competing hypotheses must correspondingly approach 0.[33] Given Confidence, this means that if we make sufficiently many observations of the systems in question, then no matter what we see, we will be no more confident in any hypothesis about how the chances work for such systems than we were initially. This would be a disaster for all scientific inquiry.[34]

One precedent for the bullet-biting response to *Two Spinners* is the reaction of Joyce (2010) to a similar argument—the inspiration for mine—against his thesis that credal imprecision is sometimes rationally *compulsory*. Joyce considers a case where you are given one of two coins chosen at random. You know the first coin is fair, but have "no information" about the bias of the second coin. He expresses sympathy for the claim that your representor should contain probability functions treating the second coin as having every bias between 0 and 1. He moreover accepts a form of representor conditionalization which implies that if you begin with such a representor, no sequence of observations should increase your lower credence that you have the first coin. Joyce suggests that the counterintuitiveness of this prescription can be defused by adopting a decision theory on which even persistent, extreme imprecision in one's credences need not show up in any distinctive patterns of *action*. For example, one might allow a process of "pragmatic sharpening" in which one proceeds *as if* one's representor were smaller than it actually is. IHH might, analogously, suggest that after your observations in *Two Spinners*, you should *act as if* you had become more confident that you have the boring spinner, without in fact becoming more confident.[35]

[32]Indeed, the hypothesis that free neutrons have lower chance 0 and upper chance 1 of decaying within any finite interval of time seems rather simple and elegant; some might find it more plausible *a priori* than the currently accepted hypothesis that their chance of decaying is approximately 0.5 per 611 seconds.

[33]The phenomenon is not limited to hypotheses that ascribe lower and upper chances of 0 or 1. The result that a long enough sequence of observations will cause one's lower and upper credences that one has the boring spinner to approach 0 and 1 respectively still holds if the set of white points on the interesting spinner has inner measure x and outer measure y and the boring spinner measure z, with $x < z < y$.

[34]As Hoëk (2021: §2) points out, ordinary scientific reasoning seems to assume the existence of real-valued chances, so it would be hard for a view on which imprecise chances deserve nonzero prior credence to be *entirely* conservative with respect to standard scientific practice. But without Permissive Gen-PP, I see no obstacle to the view that given our actual evidence, we should have negligible credence that there is any significant imprecision in the dynamical chances.

[35]Such invocations of pragmatic sharpening will exacerbate the foundational worries about credal imprecision. Behaviourists certainly won't like the idea that credal imprecision in A is compatible with being disposed to act just like someone with a particular real-valued credence. And even non-behaviourists may feel that this weakens their grip on the concept of credal imprecision.

However, it is hard to imagine what sort of normative decision theory could explain the practical requirements required by this strategy. Since Joyce was only trying to "explain away" the appeal of an initially plausible *permissibility* claim, his purposes would be served by a view on which arbitrary pragmatic sharpenings are always permissible. By contrast, the analogous strategy for "explaining away" an initially plausible 'should' judgment like mine will require a view on which some of the possible pragmatic sharpenings are *compulsory*, while others (e.g., acting as if you always had a high credence that the interesting spinner would generate the given sequence of outcomes) are impermissible. I see no basis for such a discriminating requirement.

And even if the practical claim that you should act as if you had become more confident could be integrated into a reasonable decision theory, it would still not be an adequate substitute for the claim that you should in fact become more confident. Our real-world scientific cases make this clear. Surely, for example, there are detailed observations that should really make us confident that neutron-decay involves certain fairly specific chances, not just disposed to act as if we were confident. Science is about finding out what the world is like, not just putting on a brave face!

4 Replacing the Generalized Principal Principle

I have argued that Permissive Gen-PP is false; since anything rationally obligatory is rationally permissible, it follows that Gen-PP is also false. But I don't want to go so far as to give up the original Principal Principle, so I need to say something about IHH's argument that proponents of the Principal Principle should also accept Gen-PP. Their thought is that assuming imprecise chance is possible, there clearly needs to be *some* generalization of the Gen-PP that doesn't just fall silent when applied to hypotheses according to which the chance of the relevant proposition is imprecise, and approximates the original Principal Principle when there is only a little imprecision. For example, the hypothesis that A has lower chance 0.49 and upper chance 0.51 and the hypothesis that it has real-valued chance 0.5 should presumably place similar constraints on one's attitude towards A.

I agree that some such generalization is called for. But I suggest that the following weaker principle can do the job:

> Weak Gen-PP If Cr_* and Cr^* are a rational combination of a lower and upper prior credence function, and E is admissible at t, then for any proposition A:
>
> $$\mathrm{Cr}_*(A \mid E \wedge \mathrm{Ch}_{*t}(A) = x \wedge \mathrm{Ch}^*{}_t(A) = y) \geq x$$
> $$\mathrm{Cr}^*(A \mid E \wedge \mathrm{Ch}^*{}_t(A) = x \wedge \mathrm{Ch}^*{}_t(A) = y) \leq y$$

so long as these are defined.

Like GEN-PP, WEAK GEN-PP entails the original Principal Principle: this is the special case where $x = y$.[36] When the lower and upper chance are certain, WEAK GEN-PP requires your lower and upper credences to be within the interval they span, but allows that they might be closer together than the known lower and upper chances, or even equal. WEAK GEN-PP is thus useless for the purposes of arguing for the rational permissibility of imprecise credence.

5 Symmetry

IHH do suggest a different argument that could be used to support the permissibility of imprecise credences in the spinner case, even if one accepted the impermissibility of the specific imprecise credences mandated by GEN-PP. To run the alternative argument, we'll need to further flesh out *Strange Spinner* by having God tell you that the set of white points is a *Vitali set*: one that does not overlap the result of rotating it by any rational angle, such that the union of all such rotations is the whole circle.[37] The premise of the alternative argument is that it is permissible (in this version of the case) for your beliefs about the next spin to obey both *countable additivity* and *rotational symmetry*. Countable additivity means that when you have a real-valued credence in each of countably many pairwise inconsistent propositions, your credence in their disjunction is the sum of your credences in them.[38] Rotational symmetry means that for every angle x, your lower credence that *the spinner will land on a white point* equals your lower credence that *the spinner will land on a point x clockwise from a white point*. Conditional on the set of white points being a Vitali set, you are certain there is exactly one *rational* number x such that *the spinner will land on a point x clockwise from a white point* is true. Given rotational symmetry, all of these propositions must receive the same lower and upper (conditional) credence. Since they are pairwise incompatible, this means they must all have lower credence 0. But since their disjunction gets upper credence 1 and there are only countably many of them, countable additivity says that their upper credence must be positive. Of course, if you assign each of them upper credence 1, you will be in the same state mandated by GEN-PP, which I already argued to be impermissible. But if you are willing to jettison GEN-PP, you could still obey rotation-invariance and countable additivity by assigning each of them the same upper credence strictly between 0 and 1.

[36] We assume that propositions about the imprecise chances at t are admissible at t.

[37] A Vitali set must have inner Lebesgue measure 0, but can have any outer Lebesgue measure in $[0, 1]$; its complement cannot be a Vitali set.

[38] In the imprecise setting, it is natural to derive countable additivity for credence from a requirement of countable *subadditivity* for upper credence, generalizing (A^{***}): your upper credence in a countable disjunction can never exceed the sum of your upper credences in the disjuncts.

This is not an attractive package. Insofar as rotation-invariance has any appeal independent of Gen-PP, it seems based on the idea that there is something unreasonably *arbitrary* about having credences that violate it. But choosing any particular upper credence strictly between 0 and 1 also seems awfully arbitrary. And in any case, our argument against the permissibility of the imprecise credences mandated by Gen-PP can be adapted to turn it into an argument against the permissibility of these alternative rotation-invariant imprecise credences. For any positive upper credence you might adopt, we can consider a variant of *Two Spinners* where the measure of the white points on the boring spinner is positive but less than that number. An extension of the reasoning in Equation 2 will then require your lower credence that you have this boring spinner to approach 0, whatever observations you make.

This leaves us with a choice between rejecting (the compulsoriness of) countable additivity and rejecting (the permissibility of) rotational symmetry. As it happens, I reject countable additivity for entirely independent reasons (see Arntzenius, Elga, and Hawthorne 2004; Arntzenius and Dorr 2017).[39] But (as IHH note) we can modify the case to avoid having to rely on countable additivity, by replacing our spinners with similar devices that choose points on the surface of a sphere rather than on the circumference of a circle. By the Banach-Tarski theorem, even a *finitely* additive probability function that assigns a probability to every set of points on a sphere cannot be invariant under all rotations of the sphere. So at least in these cases, I must deny that your credences should be rotation-invariant.

But this does not seem so terrible: without Gen-PP, the theoretical case for rotation-invariance does not look so strong. In the cases where considerations of symmetry seem to carry the most epistemological weight, the propositions related by a symmetry are also alike in respect of their *simplicity* (naturalness), which already speaks in favour of assigning them equal credence. For example, the judgment that one should have take the same prior credence in the propositions that a coin will land heads more often than tails and that it will land tails more often than heads is intimately connected with the fact that these two propositions are equally simple. But the proposition that the interesting spinner will land on a white point is intuitively quite a lot simpler than, e.g., the proposition that it will land on a point exactly $17\frac{9}{44}$ degrees clockwise from a white point.

[39]Rejecting countable additivity also takes care of another argument of IHH's which appeals to Ulam's theorem, according which there is no countably additive probability measure defined on the powerset of $\aleph_1$ that vanishes on singletons, and hence, if the continuum hypothesis is true, no such measure defined on the set of all propositions about where the spinner will land the first time it is spun. But a proponent of countable addivity might also respond to this argument by denying that it is permissible (assuming the conditional hypothesis is true) even to have a lower and upper credence in every proposition about where the spinner will land. This seems defensible; indeed one might argue that it is not even *possible* for creatures like us to be in such a state.

(Correspondingly, *being a white point* is a much more *natural* property than *being exactly* $17\frac{9}{44}$ *degrees clockwise from a white point*.) A theory that mandates or permits treating these propositions on a par is thus making a bold and controversial step beyond the kinds of cases that initially suggested a role for the concept of symmetry in the theory of rationality.

Could we shore up the case for rotation-invariance by switching to an example that eliminates the spinner's qualitative asymmetries (the different colours and dispositions to make the lamp come on)? The problem is that if you're just presented with a pointer attached to a featureless disc, it is unclear how you could even *entertain*, or take any other attitude to, the proposition that the pointer will end up in some specific non-measurable set of points. Once the possibility of piggybacking on God's decision to confer a qualitative distinction on one such set is removed, it is arguable that you cannot assign any lower or upper credence to any such proposition.[40] And so long as you don't, you will still be able to be rotationally invariant even if your lower and upper credences never diverge.

6 Avoiding arbitrariness

The idea that your credences should (or may) be invariant under the relevant symmetries is attractive, in part, because any specific way of *not* being symmetry-invariant seems alarmingly arbitrary. It is hard to imagine anything *principled* one could say on behalf of having a real-valued credence of $1/2$, or $1/3$, or $1/\pi$, or anything else that the lamp will come on after the first spin. The thought that we should have something principled to say in defence of every feature of our belief state is an important source of pressure towards the view that imprecise credence is sometimes mandatory. Some such thought is prominent in the work of philosophers like Joyce (2010), who are moved by the idea that the evidence we have to go on is so often so "incomplete, imprecise or equivocal" that imprecise credence is needed to respond rationally to it. If it's true (as Joyce suggests) that this vision will sometimes require one's lower and upper credences in some A conditional on some B to be 0 or 1, the argument of this paper can be adapted to become an argument against the vision. But the demand for a principled, non-arbitrary basis for each feature of one's credence function is independently suspect, since it can so readily collapse into absurd forms of scepticism. Even proponents of imprecise credence must grant that in many everyday cases, people manage to form non-extreme lower or upper credences in ways that seem reasonable, despite having no inkling of any simple first principles from which one could derive that

[40]Proponents of the inner-outer analysis will not agree, since according to them every agent must have a lower and an upper credence in every proposition. Their view makes it easy to defend the permissibility of imprecise credence, but it is itself very controversial.

those are the right credences to have.[41]

Perhaps the demand for non-arbitrariness can be understood in a more relaxed way that is content with some sort of "reflective equilibrium", rather than a derivation from first principles. This might still be thought to require extreme imprecision in our spinner case, since it is hard to muster up any non-extreme pre-theoretic intuitions about what your lower and upper credences that the lamp will come on should be, conditional on having the interesting spinner. But our argument suggests a different source of pre-theoretic judgments that could be drawn on during the progress towards reflective equilibrium, namely, judgments about which sequences of observations would and would not make it reasonable to become confident that one had this or that sort of boring spinner. We do seem to be in a position to make some piecemeal intuitive judgments about this. For example, I judged that you should be pretty confident that you have the boring spinner after the given series of 30 observations; I take this to imply that your lower credence that you have the boring spinner should then be well above 1/2. Given these judgments about how different observations would bear on the hypothesis that you have the interesting spinner, we can apply Bayes's theorem in reverse to make some progress with the question what your credences conditional on that hypothesis should look like. Of course, these judgments aren't very specific and leave quite a lot open. But in this respect, they are no different from our intuitive judgments about other cases of induction. These suggest (for example) that it is reasonable to assign higher prior credence to simpler hypotheses, but do not take us very far towards a theory about what the relevant notion of simplicity amounts to, or how the bias towards simplicity should be implemented quantitatively. The arbitrariness worries raised by propositions about unmeasurable regions thus do not, ultimately, seem importantly different from the ones that arise in connection with run-of-the-mill inductive reasoning.

7 Conclusion

I conclude that IHH's arguments for the rational permissibility of imprecise credence rest on unconvincing premises. The premise that it is permissible to satisfy the demands of Gen-PP does not seem plausible once we see what satisfying

[41] This would be denied by Sturgeon (2020), who maintains that ordinary people almost never have lower or upper credences on similar grounds to those on which many fans of imprecise credence maintain that ordinary people almost never have real-valued credences. Rinard (2013), similarly, thinks that most ascriptions of interval-valued credences should be rejected as involving a 'false precision'. But these arguments seem to turn on an anti-arbitrariness thought that cannot be endorsed by those of us who hold on to classical logic even in the face of the Sorites. (Indeed, Rinard's "false precision" argument seems pretty directly to require giving up the Law of Excluded Middle). The task of developing any analogue of precise or imprecise probability theory within non-classical logic raises challenges beyond the scope of this paper.

those demands would lead to in the spinner case. The premise that it is permissible to have rotationally symmetric lower and upper credences is subject to similar objections, and is in any case not well-motivated apart from Gen-PP. In considering these arguments, we have learnt something about the force of the anti-arbitrariness considerations which drive the other main arguments for imprecise credence in the literature. This helps clear the way for arguments against the permissibility (e.g. White 2010; Elga 2010) or possibility (e.g. Dorr, Nebel, and Zuehl 2021) of imprecise credence.

References

Arntzenius, Frank (2003), 'Some Problems for Conditionalization and Reflection', *Journal of Philosophy*, 100: 356–70. doi: 10.5840/jphil2003100729.

Arntzenius, Frank and Dorr, Cian (2017), 'Self-Locating Priors and Cosmological Measures', in Khalil Chamcham, John Barrow, Simon Saunders, and Joe Silk (eds.), *The Philosophy of Cosmology* (Cambridge: Cambridge University Press), 396–428.

Arntzenius, Frank, Elga, Adam, and Hawthorne, John (2004), 'Bayesianism, Infinite Decisions, and Binding', *Mind*, 113: 251–83. doi: 10.1093/mind/113.450.251.

Bradley, Seamus (2019), 'Imprecise Probabilities', in Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2019, Metaphysics Research Lab, Stanford University), https://plato.stanford.edu/archives/spr2019/entries/imprecise-probabilities/.

—— (2022), 'Learning by Ignoring the Most Wrong', *Kriterion*, 36/1: 9–31. doi: doi:10.1515/krt-2021-0028.

Cattaneo, Marco (2008), 'Fuzzy Probabilities based on the Likelihood Function', in D. Dubois, M. A. Lubiano, H. Prade, M. A. Gil, P. Grzegorzewski, and O. Hryniewicz (eds.), *Soft Methods for Handling Variability and Imprecision* (Dordrecht: Springer), 43–50.

Dorr, Cian, Nebel, Jacob M., and Zuehl, Jacob (2021), 'Consequences of Comparability', *Philosophical Perspectives*, 35: 70–98. doi: 10.1111/phpe.12157.

Easwaran, Kenny (2019), 'Conditional Probabilities', in Richard Pettigrew and Jonathan Weisberg (eds.), *The Open Handbook of Formal Epistemology* (PhilPapers Foundation), 131–98, https://philpapers.org/rec/PETTOH-2.

Elga, Adam (2010), 'Subjective Probabilities Should Be Sharp', *Philosophers' Imprint*, 10: 1–11. doi: 2027/spo.3521354.0010.005.

Fagin, Ronald and Halpern, Joseph Y. (1991), 'Uncertainty, Belief, and Probability', *Computational Intelligence*, 6: 160–73. doi: 10.1111/j.1467-8640.1991.tb00391.x.

Greaves, Hilary and Wallace, David (2006), 'Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility', *Mind*, 115/459: 607–32. doi: 10.1093/mind/fzl607.

Hájek, Alan (2003), 'What Conditional Probability Could Not Be', *Synthese*, 137: 273–323. doi: 10.1023/b:synt.0000004904.91112.16.

Hoëk, Daniel (2021), 'Chance and the Continuum Hypothesis', *Philosophy and Phenomenological Research*, 103/3: 639–60. doi: 10.1111/phpr.12708.

Isaacs, Yoaav, Hájek, Alan, and Hawthorne, John (2021), 'Non-Measurability, Imprecise Credences, and Imprecise Chances', *Mind*. doi: 10.1093/mind/fzab031 (forthcoming).

Jeffrey, Richard C. (1965), *The Logic of Decision* (2nd, Chicago: University of Chicago Press).

Joyce, James M. (2010), 'A Defense of Imprecise Credences in Inference and Decision Making', *Philosophical Perspectives*, 24/1: 281–323. doi: 10.1111/j.1520-8583.2010.00194.x.

Kaplan, Mark (1996), *Decision Theory as Philosophy* (Cambridge: Cambridge University Press).

Keynes, John Maynard (1921), *A Treatise on Probability* (London: Macmillan).

Koopman, B. O. (1940), 'The Bases of Probability', *Bulletin of the American Mathematical Society*, 46/10: 763–74.

Levi, Isaac (1974), 'On Indeterminate Probabilities', *Journal of Philosophy*, 71: 391–418. doi: 10.2307/2025161.

Lewis, David K. (1980), 'A Subjectivist's Guide to Objective Chance', in id., *Philosophical Papers*, i (Oxford: Oxford University Press, 1983), 83–133. From R. C. Jeffrey (ed.), *Studies in Inductive Logic and Probability*, ii (Berkeley: University of California Press, 1980), 263–93.

—— (1999), 'Why Conditionalize?', in *Papers in Metaphysics and Epistemology* (Cambridge: Cambridge University Press), 403–7.

Mahtani, Anna (2016), 'Imprecise Probabilities and Unstable Betting Behaviour', *Noûs*, 52/1: 69–87. doi: 10.1111/nous.12148.

Popper, Karl (1959), *The Logic of Scientific Discovery* (London: Hutchinson).

Popper, Karl R. (1955), 'Two Autonomous Axiom Systems for the Calculus of Probabilities', *The British Journal for the Philosophy of Science*, 6/21: 51–7. doi: 10.1093/bjps/vi.21.51.

Rinard, Susanna (2013), 'Against Radical Credal Imprecision', *Thought: A Journal of Philosophy*, 2/1: 157–65. doi: 10.1002/tht3.84.

Ruspini, Enrique H. (1987), *The Logical Foundations of Evidential Reasoning*, Technical Note 408 (Menlo Park, CA: SRI International), https://www.sri.com/wp-content/uploads/pdf/557.pdf.

Seidenfeld, Teddy and Wasserman, Larry (1993), 'Dilation for Sets of Probabilities', *The Annals of Statistics*, 21/3: 1139–54. doi: 10.1214/aos/1176349254.

Shafer, Glenn (1976), *A Mathematical Theory of Evidence* (Princeton, NJ: Princeton University Press).

Sturgeon, Scott (2020), *The Rational Mind* (Oxford: Oxford University Press).

van Fraassen, Bas (1990), 'Figures in a Probability Landscape', in Michael Dunn and Anil Gupta (eds.), *Truth or Consequences* (Dordrecht: Springer), 345–56.

Walley, Peter (1991), *Statistical Reasoning with Imprecise Probabilities* (Dordrecht: Springer).

Weatherson, Brian (2007), 'The Bayesian and the Dogmatist', *Proceedings of the Aristotelian Society*, 107/1pt2: 169–85. doi: 10.1111/j.1467-9264.2007.00217.x.

White, Roger (2010), 'Evidential Symmetry and Mushy Credence', in Tamar Szabó Gendler and John Hawthorne (eds.), *Oxford Studies in Epistemology*, iii (Oxford: Oxford University Press), 161–86.

Wilson, Nic (2001), 'Modified Upper and Lower Probabilities based on Imprecise Likelihoods', in *ISIPTA '01: Proceedings of the Second International Symposium on Imprecise Probabilities and Their Applications*, ed. Gert de Cooman, Terrence L. Fine, and Teddy Seidenfeld (Maastricht: Shaker Publishing).