

Suppose we take a standard, randomly shuffled pack of cards with no jokers, and ask what the probability is that the top card is a red picture card. We can calculate the probability of this to be $6/52 = 3/26$. And of course if, as many think, people have degrees of belief, or credences, then you—knowing only that the pack is normal and has been randomly shuffled—should have a credence of $3/26$ that the top card is a red picture card.

But many think that you can have credences in all sorts of claims, and not just claims about random events involving cards, dice or coins. In particular, classical Bayesian epistemologists think that you have a credence in every proposition that you can entertain. Thus for example, there is some number between 0 and 1 that is your credence that it will snow in London on New Year's Day 2026; and there is some number between 0 and 1 that is your credence that I have a cup of tea beside my computer as I type. But what exactly are your credences in these claims? Perhaps no particular number springs to mind. Unlike in the playing card scenario above, here there does not seem to be any obvious way to 'work out' what the probability of these events is and so arrive at the precise credence that you ought to have. Cases like these have led some to reject the classical Bayesian epistemologist's claim that people must have precise credences in every proposition that they can entertain. Instead it is claimed that people—even rational people—can have imprecise credences in at least some propositions. Hereafter I will use 'imprecise probabilism' as the name for the view that rational people can have imprecise credences.

Imprecise probabilism has some intuitive appeal. Take again the claim (which we can call 'NYD') that it will snow in London on New Year's Day 2026. It is hard to put a precise number on your credence—but there may still be something we can say about your attitude towards this proposition. For example, perhaps you think the claim is not very likely, but far from impossible, and certainly more likely than the claim (which we can call 'MIDSUMMER') that it will snow on Midsummer's day in London in 2026. We might think that your credences in these claims can be represented with ranges, rather than points. For example, perhaps your credence in NYD is the range (0.1, 0.4), and your credence in MIDSUMMER is the range (0.01, 0.05). Versions of this idea—of representing credences by ranges rather than points—can be found in numerous sources, including R. Bradley (2009), Gärdenfors and Sahlin (1982), Jeffrey (1983), Joyce (2005),

Kaplan (1996), Keynes (1921), Kyburg (1983), Levi (1974), Sturgeon (2008), van Fraassen (2006), and Walley (1991).

To explain imprecise probabilism in more depth, I must first set out the classical Bayesian view more precisely. We begin with a set (an *event space*) $\Omega = \{w_1, w_2, \dots, w_n\}$. Each w_i in Ω is a state of affairs, or possible world. We can then see a proposition (or event) A as a subset of set Ω . For example, suppose we take the proposition that a particular dice throw landed on an even number. This proposition obtains at all those possible worlds where the dice lands on 2, 4, or 6. Thus there will be a set of possible worlds where the proposition obtains. For the purposes of this topic, we assume that the proposition can be identified with that set of possible worlds at which it is true.

Now consider a set $\mathcal{F} = \{A_1, A_2, \dots, A_m\}$ of these propositions (which are themselves each sets of possible worlds). To make this set a sigma algebra (σ -algebra), the set must be closed under union, intersection, and complementation. For the set to be closed under union, it must be the case that for two propositions A_i and A_j in the set, the union ($A_i \cup A_j$) is also in the set; similarly, for the set to be closed under intersection, it must be the case that for any two propositions A_i and A_j in the set, the intersection ($A_i \cap A_j$) must also be in the set; and for the set to be closed under complementation it must be the case that for any proposition A_i in the set, the proposition $\Omega - A_i$ must also be in the set.

Finally we introduce a function p mapping \mathcal{F} to $[0, 1]$. Thus for example, if \mathcal{F} contains some proposition A , then our function p will map that proposition A to some number between 0 and 1. If the function is a probability function, then it will meet these three conditions (the probability axioms).

1. $p(A) \geq 0$, for all A in \mathcal{F} .
2. $p(\Omega) = 1$.
3. If $A_i \cap A_j = \emptyset$ then $p(A_i \cup A_j) = p(A_i) + p(A_j)$.

Classical Bayesian epistemologists claim that any rational agent has an epistemic state that can be represented by such a probability function. In contrast according to the imprecise probabilist, a rational agent may have an epistemic state that cannot be represented by a single probability function. Instead, the imprecise probabilist typically claims that a rational agent's epistemic state can be represented by a set of probability functions $P = \{p_1, p_2, \dots, p_k\}$. Thus rather than assigning each proposition in \mathcal{F} some unique number, for each proposition A in \mathcal{F} there will be some least number assigned to it by the probabilities in P (the agent's *lower envelope* of A), and some greatest number assigned to A by the probabilities P (the agent's *upper envelope* of A).

Thus the imprecise probabilist moves away from the classical Bayesian view by claiming that an agent's epistemic state is given not by a single function from propositions to numbers, but by a set of such functions. And if the agent is rational, then each of the functions in the set that represents the agent's epistemic state will be a probability function. In van Fraassen's terminology, this set of functions is the agent's *representor* (van Fraassen, 1990). On another vivid version of this account, we can see the set as a group of *avatars* (R. Bradley, 2009), each of whom has a precise credence function: these avatars collectively represent the agent's epistemic state.

This view raises some interesting problems, but before I turn to these, I will first explore in more depth the central claims of the view.

1 EXPLORING THE VIEW

How can a set of credence functions represent an agent's epistemic state? Or, to put the point another way, what must be true of a given agent for her epistemic state to be correctly represented by some specific set of credence functions?¹

The idea is that what holds across all the credence functions in the set, holds for the agent's epistemic state.² Thus for example, suppose that every credence function in the set assigns the number 0.5 to the claim that the next fair coin tossed will land heads: then it follows that the agent has a credence of precisely 0.5 in this claim. Or suppose that every credence function in the set assigns a number of no more than 0.4 to the claim *NYD*: then it follows that the agent has a credence of no more than 0.4 in this claim. Or suppose that every credence function in the set assigns a higher number to *NYD* than it does to *MIDSUMMER*: then it follows that the agent has a higher credence in the claim *NYD* than she does in *MIDSUMMER*.

On this picture, there may be some questions we can ask about the agent's epistemic state which have no answer. For example, we might wonder which of a pair of claims is given the highest credence, or whether they are given equal credence—but there may be no answer to this question if the credence functions that represent the agent's epistemic state conflict over this. Similarly, on learning that an agent has a credence of no more than 0.4 in *NYD*, we might ask what exactly the agent's credence is in this claim. But there is no answer to this question if the different credence functions that represent the agent's epistemic state assign different values to this claim. In such cases, it is natural to say that the agent's credence in a claim is a range rather than a single unique number—where the range contains all and only those numbers that are assigned to the relevant

¹ Richard Bradley argues that for any given epistemic state, there is a *unique* maximal set of such functions that represents that epistemic state (R. Bradley, 2009, p. 242).

² Or perhaps, is *determinately* true of the agent's epistemic state (Rinard, 2015).

proposition by some credence function from the set that represents the agent's epistemic state.

I turn now to consider some variations on this view, and some initial objections and clarifications.

1.1 *Variations on the view*

Here I contrast two different sorts of imprecise probabilist. All proponents of imprecise probabilism agree that agents are sometimes permitted to have imprecise credences in some propositions. They thus stand in contrast to the classical Bayesian epistemologists, according to whom rational agents have precise credences in every proposition which they can entertain. But even amongst those who accept imprecise probabilism, there is disagreement over whether imprecise credences are ever *required* by rationality.

James Joyce, for example, argues that one's degrees of belief should be no sharper than the evidence requires (Joyce, 2005): Joyce requires an agent to have an imprecise credence in a claim where the evidence for that claim does not justify a more precise credence. Thus for example consider again the claim *NYD*, that it will snow in London on New Year's Day 2026. Given that the evidence for this is (as yet) slight, an agent who had a precise credence in this claim (e.g. a credence of exactly 0.35) would be irrational. In contrast, take the claim that the next fair coin tossed will land heads. Given that the chance of this event is known to be 0.5, it is rational to have a credence of exactly 0.5 in this claim.

To clarify this view, we need to explain what determines the correct imprecise credence for an agent to have in any given situation. One possible answer to this is the *chance grounding thesis*: "one's spread of credence should cover the range of chance hypotheses left open by the evidence" (White, 2009, p. 174).³ To see what this means, let us consider a few examples. First take an agent who knows that a coin is fair, and is contemplating the claim, *HEADS*, that on the next toss the coin will land heads. Given that (s)he knows that the chance of *HEADS* is 0.5, the chance grounding thesis requires that every credence function in the set that represents the agent's epistemic state must assign 0.5 to *HEADS*—and so the agent must herself have a credence of precisely 0.5 in *HEADS*. Now suppose instead that the agent has a coin that she does not know to be fair: the chance of its landing heads (*HEADS**) is anywhere within the range (0.2, 0.8), for all she knows. Then the chance grounding thesis requires that for each value v within the range (0.2, 0.8), there must be a credence function in the set that represents the agent's epistemic state that assigns v to *HEADS**. And

³ White defines this thesis, but does not endorse it.

furthermore there must be no credence function in the set that assigns to HEADS* some value v that is outside the range $(0.2, 0.8)$.

This chance grounding thesis generates some counterintuitive results, and Joyce argues that it should be replaced with the less stringent demand that when your *only* relevant evidence is that the chance of some event is within some interval (a, b) , then your spread of credence ought to cover this range (Joyce, 2010, p. 289). So for example suppose that in the case above, you know not only that the chance of the coin's landing heads is within the range $(0.2, 0.8)$, but also that the coin was selected at random from a bag which contained a variety of coins with complementary biases: i.e. for each coin in the bag that has a chance v of landing heads, the bag also contained exactly one coin with a chance $1 - v$ of landing heads. In this case, because you have this extra piece of evidence, your "spread of credence" in HEADS is not required to cover the whole range $(0.2, 0.8)$, and a credence of precisely 0.5, say, is permitted. However if you know only that the chance of the coin's landing heads is within the range $(0.2, 0.8)$, then your spread of credence in HEADS is required to cover the whole range $(0.2, 0.8)$.

Now we turn to consider imprecise probabilists who permit, but never require agents to have imprecise credences. For these theorists, an agent is free to have a credence of precisely 0.35 in the claim NYD (that it will snow in London on New Year's Day 2026). To these theorists, we might ask whether there are any rational constraints on an agent's epistemic state, bar the requirement that their state should be represented by some maximal set of credence functions that obey the probability axioms. Such a theorist might require that any rational agent's epistemic state will conform to the *principal principle*—i.e. that the agent's credence in any claim P conditional on the chance of P being some value v , is v (Lewis, 1980). From this, it follows that in the case where an agent is contemplating the claim (HEADS) that on its next toss a coin known to be fair will land heads, the agent's credence in HEADS must be 0.5. But what constraint is placed on the agent in the case where (s)he is contemplating the claim (HEADS*) that on its next toss a coin known to have a chance within the range $(0.2, 0.8)$ will land heads? The principal principle here requires that the agent's credence should not exceed the range $(0.2, 0.8)$, but nothing seems to require that the agent's credence should occupy this entire range.

Having explored this variation in the views of imprecise probabilists, I turn now to contrast the account with an alternative view.

1.2 Dempster-Shafer Theory

An alternative approach to modelling our epistemic state involves *belief functions* (Dempster, 1967, 1968; Shafer, 1976). To illustrate this view, we

can again take the proposition (NYD) that it will snow in London on New Year's Day in 2026, and suppose that my belief-function assigns a value of 0.6 to this claim: we represent this by writing $Bel(\text{NYD}) = 0.6$. If my belief function was a probability function, then it would follow that the value assigned to the negation of NYD (i.e. to not-NYD) would be 0.4. However a belief function need not be a probability function, and it might assign any value less than or equal to 0.4 to not-NYD. Thus for example, it might assign a value of 0 to not-NYD. This is despite the fact that the value assigned to the tautology (either NYD or not-NYD) must be 1.

More generally, on this view the value assigned to the disjunction of two disjoint propositions A and B , $Bel(A \cup B)$, need not equal the sum of $Bel(A)$ and $Bel(B)$. The requirement is only that the value assigned to the disjunction must be at least as great as the sum of the values assigned to the disjuncts. Thus the belief function is not a probability function, as the third probability axiom (countable additivity) does not apply.

One way to interpret the idea of a belief function, is as a measure of the weight of evidence for each proposition. Thus consider again my belief function that assigns a value of 0.6 to NYD. We can suppose that I have asked a friend whether it will snow in London on New Year's Day 2016, and (s)he assures me that it will. I consider this friend to be reliable in 60% of cases of this sort, and this explains why my belief function assigns a value of 0.6 to this claim. If we suppose that this is all the relevant evidence that I have, then my belief function assigns a value of 0 to not-NYD simply because I have no evidence to support not-NYD. In cases where I have evidence from two different sources (e.g. in a case where I make another friend who also gives me his or her opinion on NYD), then the belief functions that result from these different bodies of evidence need to be combined, and Dempster and others have explored the question of how this combination should be carried out (Dempster, 1967).

In common with imprecise probabilism—and in apparent contrast with classical Bayesianism—this theory has resources designed to model severe uncertainty. To see this, suppose that a coin is about to be tossed, and that you have no information whatsoever about whether the coin is fair or how it might be biased. On the classical Bayesian view, in spite of your severe uncertainty, you will nevertheless have a precise probability that the coin will land head-side-up. This strikes many as counterintuitive. Advocates of both imprecise probabilism and Dempster-Shafer theory take their theories to improve on classical Bayesianism here. According to imprecise probabilism, in the case where you have no information about the bias of the coin, a rational agent may—and on some versions of the theory, must—have a credal range of $(0, 1)$ rather than a precise credence of 0.5. And according to Dempster-Shafer theory, in a case where you have no information about the bias of the coin, you have no evidence in favour

of heads, and no evidence in favour of tails, and so your belief function will assign a value of 0 to both HEADS and TAILS.

For more on the Dempster-Shafer theory, and how it differs from both classic Bayesianism and imprecise probabilism, see Halpern (2003) and Yager and Liu (2008).

1.3 Scoring Rules

I turn now to an issue for those theorists who want to apply the idea of accuracy *scoring rules* in the context of imprecise probabilism. I begin by outlining a standard proposal for measuring the (in)accuracy of a credence function, and I explain how this sort of scoring rule has been used to construct an argument for probabilism. I then gesture towards some of the challenges that arise when we consider these measures of accuracy in the context of imprecise probabilism.

Let's begin then with the classical Bayesian picture, according to which a rational agent's epistemic state is represented with a single precise credence function. In this context a variety of scoring rules have been proposed for measuring a credence function's (in)accuracy at a given world. One popular such rule is the *Brier score* (Brier, 1950) which I outline here. First we set the truth-value of a proposition at a world to 1 if the proposition is true there, and 0 if it is false. Now we can measure the "distance" between the truth-value of the proposition at a world and the credence assigned to it, by taking the difference between the two and squaring the result. To illustrate this, suppose that you have a credence of 0.8 in the proposition that the world's population is over 7 billion in 2016. In the actual world, this proposition is true, and so has a truth-value of 1. Thus we measure the distance between the credence assigned to this proposition and its truth-value in the actual world as follows: take the truth value of the proposition (1), deduct the value assigned to it by the credence function (0.8), and then square the result (giving 0.04). We get the inaccuracy score for an entire credence function at a world by calculating this distance for each proposition that is assigned a value by the credence function, and summing the lot.

The Brier score is just one suggestion for measuring inaccuracy, and others have been proposed, along with various claims about conditions that any scoring rule ought to fulfil. One such requirement is that a scoring rule ought to be *proper*, which can be defined as follows: any agent with a rationally permissible credence function (i.e. one that obeys the probability axioms), will score her own credence function to be no more inaccurate than every other credence function, if the scoring rule that she uses is proper. The Brier score is one example of a scoring rule that meets this requirement.

Scoring rules of this sort have been used to argue for probabilism—i.e. for the claim that a rational agent's credence function obeys the probability axioms. The argument works by showing that for any credence function Cr that does not obey the probability axioms, there is an alternative credence function Cr^* which does obey the probability axioms and which dominates Cr in the following sense: the inaccuracy of Cr is at least as great as the inaccuracy of Cr^* at every world, and at some world the inaccuracy of Cr is greater than the inaccuracy of Cr^* . Thus, the argument goes, it would be irrational to have a credence function such as Cr which does not obey the probability axioms, when an alternative credence function Cr^* is available. Arguments of this sort can be constructed using any scoring rule provided that it meets certain requirements—including the requirement that it be proper (Joyce, 1998). Arguments from accuracy for a variety of other epistemic principles have also been proposed, including an argument for the principal principle (Pettigrew, 2013), and conditionalization (Greaves & Wallace, 2006).

We can now consider how these issues are affected by a switch from precise to imprecise probabilities. If an agent has an imprecise credence function, then how should the inaccuracy of her credence function be measured? We can see at once that the original measures of inaccuracy cannot be straightforwardly carried across—for where an agent's credence in some proposition is imprecise, we have no single number which measures that agent's credence, and so cannot make sense of the idea of deducting the agent's credence in a given proposition from its truth-value at some world. Thus a new way of measuring inaccuracy is needed.

There is not yet any consensus as to what this new way of measuring inaccuracy would be like. Some authors have proposed requirements that any way of measuring the inaccuracy of an imprecise credence function would need to meet, and some have uncovered difficulties for the project. Seidenfeld, Schervish, and Kadane argue that there is no strictly proper scoring rule for imprecise probabilities. See Seidenfeld, Schervish, and Kadane (2012) and Mayo-Wilson and Wheeler (2016) for further discussion on this issue. Schoenfield (2017) argues that if the new accuracy scoring rule meets certain conditions, then the claim that accuracy is all that matters is incompatible with the claim that imprecise probabilities are sometimes rationally required—or even permitted. Thus challenges await those who wish to endorse both imprecise probabilism and accuracy arguments.

Having explored the account of imprecise probabilities, I turn now to some of the most discussed objections and problems for the account. I divide these into two categories: learning and deciding.

2 LEARNING

On the classic Bayesian picture, an agent's epistemic state is represented by a single credence function. If the agent is rational, then she will update (only) by conditionalization. Thus for example suppose that an agent is about to run an experiment at the end of which she will have learnt (just) either E or not- E . At the start of the experiment (at t_0) let's suppose that the agent has a credence of 0.2 in E , and a credence of 0.5 in some hypothesis H . Furthermore, the agent has a conditional credence of 0.9 in H given E : in other words, if we let Cr_0 name the agent's credence at t_0 , then $Cr_0(H | E) = Cr_0(H \cap E) / Cr_0(E) = 0.9$. Now suppose that the experiment runs, and at t_1 the agent discovers E . The agent's new t_1 credence function (Cr_1) ought rationally to be her old t_0 credence function (Cr_0) conditionalized on the new evidence that she has gained, E . Thus her new credence in H ought to be her old conditional credence in H given E : $Cr_1(H) = Cr_0(H | E) = 0.9$.

For the proponent of imprecise probabilities, an agent's epistemic state is represented by a set of credence functions. How will a rational agent adjust her epistemic state in the light of evidence on this account? The idea standardly endorsed by imprecise probabilists is that each credence function in the set will be adjusted in the usual way by conditionalization, and the agent's new, post-evidence epistemic state can be represented by this adjusted set of credence functions. Thus for example, to return to our experiment case above, suppose that every credence function in the set that represents the agent's epistemic state at t_0 assigns a number within the range $(0.4, 0.6)$ to H —and every number within this range is assigned to H by some credence function in the set. And suppose furthermore that for each of these credence functions, the conditional credence assigned to H given E is within the range $(0.85, 0.95)$ —and every number within this range is the conditional credence assigned to H given E by some credence function within the set. Then at t_1 , when the agent has learnt (just) E , the agent's epistemic state will be represented by the original set of credence functions each conditionalized on E , and thus the agent's new credence in H will be given by the range $(0.85, 0.95)$. I will now turn to two problems—both related to learning—for the proponent of imprecise probabilities.

2.1 *Belief Inertia*

Let us consider a scenario in which you have just selected a coin from a bag, knowing only that the bag contains various coins some of which may be biased to various unspecified degrees. You are going to toss the coin 25 times, and before you begin tossing the coin (a time we can call t_0) you

contemplate claim HEADS_{25} —the claim that the coin will land heads on its 25th toss. According to any proponent of imprecise probabilities, you are permitted to have an imprecise credence in this claim. Now we can consider what will happen to your credence in HEADS_{25} if you toss the coin a few times, and it lands heads each time. Let HEADS_1 be the claim that the coin lands heads on the first toss, HEADS_2 be the claim that the coin lands heads on its second toss, and so on. Intuitively, your credence in HEADS_{25} ought to increase on learning HEADS_1 , and increase even more on learning $(\text{HEADS}_1 \cap \text{HEADS}_2)$, and so on.

For a certain sort of proponent of imprecise probabilism, this scenario is problematic. In particular, consider the sort of imprecise probabilist who claims that an agent's epistemic state should conform to the chance grounding thesis.⁴ On this view, all and only those credence functions which are compatible with the known chances must be included in the set that represents the agent's epistemic state. In the scenario that we are considering, at t_0 you can rule out very few chance hypotheses: for all you know, the chance of HEADS_{25} may be any number strictly between 0 and 1. Thus at t_0 your credence in HEADS ought rationally to be the range $(0, 1)$. What happens if you toss the coin once and it lands heads—i.e. if you learn HEADS_1 ? For any number n within the range $(0, 1)$, you have not learnt that the chance of HEADS_{25} is not n . For example, you have not learnt that the chance of HEADS_{25} is not 0.0001. Thus your new credence in HEADS_{25} , after learning HEADS_1 , ought still to be the range $(0, 1)$. What happens if you toss the coin again, and it again lands heads—i.e. in addition to HEADS_1 , you also learn HEADS_2 ? You cannot then rule out any additional chance hypotheses. For example, it may still be the case, for all you know, that the chance of HEADS_{25} is 0.0001. Thus your credence in HEADS_{25} after learning both HEADS_1 and HEADS_2 remains the range $(0, 1)$. This pattern continues: even if you toss the coin 24 times and it lands heads on each toss, your credence in HEADS_{25} should still remain fixed at $(0, 1)$. In this sense, your epistemic state exhibits inertia in the face of evidence. That your epistemic state should rationally exhibit this inertia is very counterintuitive: surely as you toss the coin and it lands heads repeatedly, your credence in HEADS_{25} ought to increase?

To put the point vividly, we can imagine the credence functions that represent your epistemic state as a group of avatars. The avatars at t_0 will assign various precise credences to HEADS_{25} : for every number in the range $(0, 1)$, there will be some avatar who assigns that value to HEADS_{25} . On learning HEADS_1 , each avatar ought to update accordingly by conditionalizing. Take an avatar who had a credence of 0.0001 in HEADS_{25} .

⁴ A similar problem applies to Joyce's adjusted version of this principle mentioned earlier.

It may be⁵ that this avatar's conditional credence in HEADS₂₅ given HEADS₁ is higher than her unconditional credence in HEADS₂₅, in which case this avatar will increase her credence in HEADS₂₅ on learning HEADS₁. But there will be some avatar (perhaps an avatar whose unconditional credence in HEADS₂₅ was even lower than 0.0001) whose credence in HEADS₂₅ conditional on HEADS₁ is 0.0001. Thus even after learning HEADS₁, there will still be, in the set representing your epistemic state, an avatar whose credence in HEADS₂₅ is 0.0001. Similarly, even if you learn the conjunction of the claims HEADS₁ through HEADS₂₄, there will still be an avatar in the set representing your epistemic state whose credence in HEADS₁ is 0.0001. Thus your credence in HEADS₂₅ will not shift from the range (0, 1) no matter how much evidence you amass in favour of HEADS₂₅.

This looks like a problem—at least for those imprecise probabilists who accept the chance grounding thesis, or something close to it. For some of the responses available, see R. Bradley (2017), Joyce (2010), Rinard (2013), and Vallinder (2018).

2.2 Dilation

Here we turn to another problem for the proponent of imprecise probabilism. The phenomenon I discuss here was first noted by early statisticians of imprecise probabilism Walley (1991) and Seidenfeld and Wasserman (1993), and has recently been prominently discussed by White (2009). Take some claim P , that you have no evidence whatsoever for or against, so that your credence at t_0 in P is the range $[0, 1]$. Suppose that I know whether P is true, and I take a fair coin and paint the heads side over. I write " P " on this heads side iff P is true, and "not P " on the heads side iff P is not true. I similarly paint over the tails side of the coin, and write on this side whichever claim (out of " P " and "not P ") is false. You know that I have done this. I then toss the coin before your eyes. Your credence before it lands (i.e. at t_0) that it will land head-side up (HEADS), is 0.5. Then at t_1 you see it land, with the " P "-side up. What then at t_1 is your credence in P and what is your credence in HEADS?

At t_1 you have learnt that the coin has landed " P "-side up. Thus if P is true, then HEADS is also true (i.e. it must have landed heads)—for if P is true then " P " has been painted onto the heads side of the coin, and so given that it has landed " P "-side up it has also landed heads. Furthermore, if HEADS is true, then P is also true—for if it has landed heads then given that it has landed " P "-side up, " P " must have been painted onto the heads side of the coin, which will have happened only if P is true. Thus at t_1 you

⁵ Though it need not be: perhaps some avatars will stubbornly refuse to adjust their credence in HEADS₂₅ from 0.0001. We might try to avoid this problem by excluding such agents (Halpern, 2003), though this will not solve the problem discussed in the main text.

can be certain that P is true iff HEADS is true. Thus at t_1 you must have the same credence in P as you have in HEADS. Given that at t_0 your credence in HEADS is 0.5, and your credence in P is the range $[0, 1]$, how will your credence adjust between t_0 and t_1 ? Will your credence in HEADS become the range $[0, 1]$? Or will your credence in P become precisely 0.5? Both options seem counterintuitive.⁶ It seems implausible that your credence in HEADS should “dilate” to the range $[0, 1]$: surely (by the principal principle) your credence that a fair coin has landed heads ought to be 0.5, unless you have some evidence as to how it has landed. And knowing that it landed on the “ P ”-side does not seem to give you any evidence as to whether it has landed heads or tails. And it also seems implausible that your credence in P should sharpen to the number 0.5 (White, 2009), for after all you knew even at t_0 that the coin would either land “ P ”-side up, or “ P ”-side down, and we cannot say that learning either of these pieces of information would force your credence in P to become precisely 0.5 without violating van Fraassen’s reflection principle (van Fraassen, 1984).

One popular response made by the imprecise probabilist, is to accept that at t_1 your credence in HEADS ought to dilate to $[0, 1]$.⁷ Here are two things that might be said in defence of this position.

- It seems as though learning that the coin landed “ P ”-side up gives you no evidence as to whether it has landed head-side up. But this would not follow if P was a claim that you knew something about. Suppose as a contrast case, then, that P is the claim that you have just won the lottery—a claim in which you have a very low credence indeed. On hearing that I (who know the outcome) am painting the true claim (out of “ P ” and “not- P ”) on the heads side, and the false claim on the tails side, you will be almost certain that I am painting “not- P ” on the heads side, and “ P ” on the tails side. Your credence at t_0 in HEADS is 0.5, but when at t_1 you learn that the coin has landed “ P ”-side up, you will be almost certain that HEADS is false. Thus where you have some evidence concerning P , it is natural to suppose that learning that the coin has landed “ P ”-side up will alter your credence in HEADS (see Sturgeon, 2010, Joyce, 2010).

What about in the case where P is a claim about which you have no evidence? In this case, it is tempting to suppose that learning that the coin has landed “ P ”-side up gives you no reason to adjust your credence in HEADS. But the situation is more complicated than

6 A further option would be for both your credence in HEADS and your credence in P to adjust, but this is no more appealing than the alternatives.

7 As White acknowledges, some statisticians and philosophers (such as Walley, 1991, and Seidenfeld and Wasserman, 1993) had noted this result and “taken it in their stride” (White, 2009, p. 177).

this suggests. Consider again your epistemic state as a set of avatars. For every number in the range $[0, 1]$, there is some avatar in the set that represents your epistemic state that assigns this number to P . Each such avatar, on learning that the coin has landed “ P ”-side up, will adjust her credence in HEADS accordingly.⁸ For example, the avatar whose credence in P is 0.2 will adjust her credence in HEADS downwards; and the avatar whose credence in P is 0.8 will adjust her credence in HEADS upwards. More generally after conditionalizing on the claim that the coin has landed “ P ”-side up, for every number in the range $[0, 1]$, there will be an avatar who assigns that number to HEADS. We can see then that it is not that learning that the coin has landed “ P ”-side up gives you no evidence relevant to HEADS, but rather that you are just very uncertain as to in what direction the evidence you have received should pull you, and how far. Thus your credence in HEADS is infected with the imprecision that you assigned to P , and your credence in HEADS dilates to the range $[0, 1]$ (Joyce, 2010).

- It is tempting to object that it is counterintuitive for an increase in evidence to leave your credence function more imprecise than it was before. However it is not obvious that your credence function is more imprecise at t_1 than it was at t_0 . To see this, consider that at t_0 though your credence in HEADS was precise, your conditional credence in HEADS given that the coin lands “ P ”-side up was imprecise. Thus there was imprecision in your credence function even at t_1 : this just was not obvious when we focused only on your unconditional credence in HEADS (R. Bradley, 2017).

Further discussion of the problem of dilation can be found in R. Bradley (2017), S. Bradley and Steele (2014), Dodd (2013), Joyce (2010) and Pederson and Wheeler (2014).

3 DECISION-MAKING

On the classic Bayesian picture, a rational agent has a precise credence function assigning some number between 0 and 1 to each proposition, and also a precise utility function assigning some number to each possible outcome representing in some sense how much the agent values each outcome. When faced with a decision problem—i.e. a choice between different actions—on the classic picture the agent must choose an action that has maximum *expected utility*. We can calculate the expected utility of any given action for the agent as follows: for every possible outcome,

⁸ Those avatars whose credence at t_0 in P is 0.5 need make no adjustment.

	Milk at home (s_1)	No milk at home (s_2)
	$Cr(s_1) = 0.5$	$Cr(s_2) = 0.5$
STOP FOR MILK	9	9
DON'T STOP	10	5

Table 1: A decision problem

we multiply the agent's credence that the outcome will obtain should she perform the action under consideration, by the utility of that outcome—and then we sum the lot.⁹

Here is an example to illustrate this. Sometimes on the way home from work, I stop to buy a pint of milk, which means that I take a bit longer to get home, but it is certainly better than getting home and finding that there is no milk in the house. Suppose that on this occasion, my credence that there is milk in the house already is 0.5. Table 1 represents my assessment of the possible outcomes.

We can now calculate the expected utility of each available action. The expected utility of stopping to buy milk is $(0.5)(9) + (0.5)(9) = 9$, whereas the expected utility of not stopping to buy milk is $(0.5)(10) + (0.5)(5) = 7.5$. On the classic decision rule “maximise expected utility”, I ought to stop to buy milk, because this is the action with the highest expected utility.

The maximise expected utility rule works on the assumption that for every relevant state of the world, the rational agent has a precise credence that that state of the world obtains. But proponents of imprecise probabilities deny this, and so cannot accept this rule. What alternative rule should they put in its place? According to the proponent of imprecise probabilities, what requirements does rationality place on an agent's choice of action? Many different answers have been proposed, and I will briefly outline two of these answers.

PERMISSIVE CHOICE RULES Recall that we can see an agent's epistemic state as represented by a set of avatars, each with a precise credence function. Thus faced with any decision problem, each avatar will have a view as to which action—or actions—will maximise expected utility.¹⁰ According

⁹ This is a rough and ready sketch of Savage's account (Savage, 1954). Modifications have been made to that account (e.g. in Jeffrey, 1965) but here I will stick to straightforward examples so that the modifications should not be relevant.

¹⁰ Here I assume that the agent has a precise utility function which feeds into each avatar's calculation. This of course is also up for debate, and some argue that just as a rational agent can have an imprecise credence function, so (s)he can have an imprecise utility function. I do not discuss this further here however.

to the permissive choice rules,¹¹ the agent may rationally perform any action provided that at least one of her avatars recommends that action.

To illustrate this, suppose that an agent's credence that it will rain tomorrow is the range $(0.4, 0.8)$. Thus for every number in this range, there is some avatar who assigns that number to the claim that it will rain tomorrow. Suppose then that the agent is offered the following bet: she is to pay out £5, and will get £10 back iff it rains tomorrow. The agent has to choose whether to accept the bet, or reject it. We can assume that the agent values only money, and values it linearly. Some of her avatars would recommend accepting the bet (those whose credence that it will rain is greater than 0.5), some recommend rejecting it (those whose credence that it will rain is less than 0.5), and some rate the expected utility of accepting it equal to the expected utility of rejecting it (those whose credence that it will rain is 0.5). Thus according to the permissive choice rules, the agent is free to either accept or reject the bet: both actions are permissible. This rule—together with some variations—is discussed under the name 'Caprice' by Weatherson (1998).

MAXIMIN The rule maximin works as follows. Where an agent has an imprecise probability function, we can see her epistemic state as represented by a set of precise functions, or avatars. When considering a possible action, there is an expected utility for that action relative to each precise probability function in the agent's set. Amongst these expected utilities for the action, one will be the lowest—and so each action has a minimum expected utility. According to maximin, when faced with a choice, a rational agent will carry out whichever action has the maximum minimum expected utility.

To illustrate this, take again our agent whose credence that it will rain tomorrow is the range $(0.4, 0.8)$: for every number in this range, there is some avatar who assigns that number to the claim that it will rain tomorrow. Suppose then that the agent is offered the following bet: she is to pay out £5, and will get £10 back iff it rains tomorrow. Each avatar calculates the expected utility of each possible action—i.e. the action of accepting the bet and the action of rejecting the bet. The avatar who assigns the lowest expected utility to accepting the bet is the avatar whose credence that it will rain tomorrow is 0.4: assuming again that the agent values only money and that linearly, we can represent the expected utility of accepting the bet from the perspective of this avatar as $-5 + (0.4)(10) = -1$. Thus the minimum expected utility of accepting the bet is -1 . Now we can calculate the minimum utility of rejecting the bet. All avatars assign the same expected utility to this action—namely 0. Thus the minimum expected utility of rejecting the bet is 0. A rational agent will choose from amongst

¹¹ This is Elga's (2010) term.

those actions with the highest minimum expected utility—and as rejecting the bet has a higher minimum expected utility (0) than accepting the bet (−1), the agent if rational will reject the bet.

Variations on this rule have been developed by Gärdenfors and Sahlin (1982), Gilboa and Schmeidler (1989), and others. An analogous maximax rule has been developed by Satia and Lave (1973). Many further rules have been proposed, including those by Arrow and Hurwicz (1972) and Ellsberg (1961). See Troffaes (2007) for a discussion and comparison of some of these rules.

3.1 *Applying These Rules*

In some scenarios, some of the alternative rules developed by imprecise probabilists seem to work better than the classical Bayesian’s rule maximise expected utility. Here is a famous case—the Ellsberg paradox—in which this holds (Ellsberg, 1961).

You have an urn before you, which contains 150 balls. 50 are black, and the other 100 are some mixture of red and yellow—but you have no further information as to what the proportions of red and yellow balls are. For all you know, there may be 100 red balls and no yellow balls, or 100 yellow balls and no red balls, or any mixture between these two extremes. Now a ball will shortly be selected at random from the urn, and you have the chance to bet on what colour the ball will be. You can either say ‘black’, in which case you’ll win £100 if it is black, and nothing otherwise; or you can say ‘red’, in which case you’ll win £100 if it is red, and nothing otherwise (Table 2).

	Black (B)	Red (R)	Yellow (Y)
BET BLACK	£100	£0	£0
BET RED	£0	£100	£0

Table 2: The first scenario in the Ellsberg paradox

Now suppose instead that you have the option of saying ‘black or yellow’, in which case you’ll win £100 if the ball is either black or yellow, and nothing otherwise; or you can say ‘red or yellow’, in which case you’ll win £100 if the ball is either red or yellow, and nothing otherwise (Table 3).

Typically people choose to say ‘black’ in the first scenario, but ‘red or yellow’ in the second. Furthermore, many apparently rational people exhibit this betting pattern.¹² The problem is that if we assume that a rational agent has precise probabilities and utilities, and chooses only

¹² See Voorhoeve, Binmore, Stefansson, and Stewart (2016) for an analysis and discussion of the prevalence of this betting pattern.

	Black (<i>B</i>)	Red (<i>R</i>)	Yellow (<i>Y</i>)
BET BLACK OR YELLOW	£100	£0	£100
BET RED OR YELLOW	£0	£100	£100

Table 3: The second scenario in the Ellsberg paradox

between those actions that maximise expected utility, then a rational agent cannot exhibit this betting pattern. To see this, let’s suppose that some agent who exhibits this betting pattern has precise probabilities, and is maximising expected utility. We let the agent’s credence in *B*, *R* and *Y* be given by $Cr(B)$, $Cr(R)$ and $Cr(Y)$ respectively, and we let the utility of winning £100 be given by u_1 and the utility of winning £0 be given by u_2 . Then—given that our agent chooses ‘black’ over ‘red’ in the first scenario, it follows that

$$Cr(B) \cdot u_1 + Cr(R) \cdot u_2 + Cr(Y) \cdot u_2 > Cr(B) \cdot u_2 + Cr(R) \cdot u_1 + Cr(Y) \cdot u_2,$$

and so that

$$Cr(B) \cdot u_1 + Cr(R) \cdot u_2 > Cr(B) \cdot u_2 + Cr(R) \cdot u_1.$$

But then the agent chooses ‘red or yellow’ over ‘black or yellow’ in the second scenario, and so it follows that

$$Cr(B) \cdot u_1 + Cr(R) \cdot u_2 + Cr(Y) \cdot u_1 < Cr(B) \cdot u_2 + Cr(R) \cdot u_1 + Cr(Y) \cdot u_1,$$

and so that

$$Cr(B) \cdot u_1 + Cr(R) \cdot u_2 < Cr(B) \cdot u_2 + Cr(R) \cdot u_1.$$

This contradicts our earlier result. Thus no agent exhibiting this betting pattern can have only precise probabilities and utilities and be guided by the rule maximise expected utility.

What alternative rule might be guiding the agent’s behaviour in Ellsberg’s scenario? Several of the rules formulated by proponents of imprecise probabilities can explain the agent’s behaviour, and so Ellsberg’s scenario can be used to argue both for (some of) the alternative rules, and for the claim that rational agents can have imprecise probabilities. To illustrate how some of these rules might handle Ellsberg’s scenario, I will run through Ellsberg’s own solution to the problem.

In Ellsberg’s terminology, a situation can be “ambiguous” for an agent. In an ambiguous situation, more than one probability distribution seems reasonable to the agent. We can gather these probability distributions into a set $P = \{p_1, p_2, \dots, p_n\}$: these are the distributions that the agent’s information “does not permit him to rule out” (Ellsberg, 1961, p. 661). The

agent assigns weights to each of these reasonable distributions, and arrives at a composite “estimated” distribution p_i where p_i is a member of P . The *estimated pay-off* A_{est} of a given action A is the expected utility of the action calculated using p_i (Ellsberg, 1961, p. 661). But when faced with a choice of actions, the rational agent may be guided not just by the expected pay-off of each action, calculated in terms of p_i . The agent may also take into account the lowest expected utility of each action as calculated using any member of P . We let A_{min} denote the minimum expected utility of action A as calculated using any member of P , and we let x denote the agent’s degree of confidence in p_i (the “estimated” distribution). Then the *index* of an action A is given by $x \cdot A_{est} + (1 - x) \cdot A_{min}$. Ellsberg’s rule for action, then, is as follows: choose the action with the highest index.

In Ellsberg’s scenario, the agent is in an ambiguous situation: the agent can be certain that the probability that a ball randomly drawn from the urn will be red is $1/3$, but the agent cannot be certain of the probability of the ball’s being yellow or black, because (s)he does not know the proportion of yellow and black balls in the urn. There are a range of probability distributions that seem reasonable to the agent: for every number r between 0 and $2/3$, there is a reasonable probability distribution under which the probability of R is r , the probability of Y is $2/3 - r$, and the probability of B is $1/3$. Let us assume for simplicity that the agent assigns weight evenly across these reasonable probability distributions. Thus on the composite “estimated” distribution, the probability of R is $1/3$, the probability of Y is $1/3$, and the probability of B is $1/3$. Thus the expected payoff of saying ‘black’ in the first scenario ($1/3 \cdot u_1 + 2/3 \cdot u_2$) is the same as the expected payoff of saying ‘red’ in that scenario, and the expected payoff of saying ‘black or yellow’ in the second scenario ($2/3 \cdot u_1 + 1/3 \cdot u_2$) is the same as the expected payoff of saying ‘red or yellow’ in that scenario.

However a rational agent need not be guided merely by the estimated payoff of each action, but also by the lowest expected utility of each action. For the action of saying ‘red’ in the first scenario, the lowest expected utility is that given by the probability distribution according to which the probability of R is 0, the probability of Y is $2/3$, and the probability of B is $1/3$: according to this distribution, the expected utility of saying ‘red’ is 0. In contrast, according to every distribution the expected utility of saying ‘black’ is $1/3$, and so of course the lowest expected utility of saying ‘black’ is $1/3$. The ‘index’ of some action A is given by $x \cdot A_{est} + (1 - x) \cdot A_{min}$, where x is the agent’s level of confidence in the ‘estimated distribution’. Thus the index of saying ‘red’ is $1/3 \cdot x + 0 \cdot (1 - x)$, and the index of saying ‘black’ is $1/3 \cdot x + 1/3 \cdot (1 - x)$. Thus whenever the agent is less than perfectly confident in the estimated distribution—which a rational agent may well be—the value x will be less than 1, and the index of saying ‘black’ will be greater than the index of saying ‘red’. Thus any agent for whom x is less

than 1 will say 'black' rather than 'red' in the first scenario. In the second scenario, however, the very same agents will choose to say 'red or yellow' rather than 'black or yellow'. For it works out that the expected payoff of both of these actions is $2/3$, but the lowest expected utility of saying 'black or yellow' ($1/3$) is lower than the lowest expected utility of saying 'red or yellow' ($2/3$), and so saying 'black or yellow' has a lower index than saying 'red or yellow'.

In short, an agent for whom x is less than 1 is *ambiguity averse*: all else being equal, the agent prefers actions where (s)he knows the chances of the relevant outcomes over actions where (s)he merely estimates those outcomes. In the first scenario, if the agent says 'black' then (s)he will know the chance of winning £100, whereas if she says 'red' then the chance of winning will be unknown. In contrast, in the second scenario, if the agent says 'red or yellow' then (s)he will know the chance of winning £100, whereas if she says 'black or yellow', the chance of winning will be unknown. Thus the betting pattern that is typically displayed in Ellsberg's scenario is permissible.

Here the imprecise probabilist seems to have an advantage over the precise probabilist. The precise probabilist seems forced to claim—counterintuitively—that the typical betting pattern in Ellsberg's scenario is irrational, whereas the imprecise probabilist can account for this betting pattern well.

I turn now to the problem of sequential decision problems, which seem to pose a problem for the imprecise probabilist.

3.2 Sequential Decision Problems

Here is a problem posed by Elga (2010).¹³ According to the imprecise probabilist, a rational agent may have a credence of, say, $[0.1, 0.8]$ in some claim H . Now consider the following two bets:

Bet A: If H is true, then you lose £10; otherwise you win £15.

Bet B: If H is false, then you lose £10; otherwise you win £15.

These bets are offered sequentially: first Bet A is offered to the agent, and then Bet B. The agent knows that she will be offered both bets, and has the option of taking both, rejecting both, or taking either one. Intuitively, it would be irrational for an agent to reject both bets, because rejecting both bets leaves the agent with nothing, whereas accepting both bets leaves the agent with a sure £5. Surely then a rational agent would not reject both?

¹³ The problems that the imprecise probabilist faces over sequential decision problems are widely discussed in the literature from economics, and a puzzle related to Elga's can be found in Hammond (1988).

The challenge that Elga poses to the imprecise probabilist is to put forward a plausible decision rule that entails that a rational agent in this scenario will not reject both bets. Various attempts have been made to meet this challenge.

It seems at first as though the permissive choice rules will not do. To see why, consider that if the agent is presented with just Bet A, there will be avatars who recommend rejection, so it follows that the agent is rationally permitted to reject Bet A. But then when presented with Bet B, there will similarly be avatars (different avatars) who recommend that this bet is rejected. So it follows that the agent is rationally permitted to reject Bet B. Thus it seems that the permissive choice rules would permit the agent to reject both bets, and so this rule cannot be used to meet Elga's challenge. However defenders of this rule may claim either that a sequence of actions is permitted only when that sequence is recommended by a single avatar, or else challenge Elga on his assumption that accepting each bet is a separate action, rather than parts of a single action (Weatherson, 2003; Williams, 2014).

Similarly, it may seem that maximin, Ellsberg's rule, and others will be unable to handle Elga's scenario, for many of these rules would permit a rational agent to reject both bets if offered on separate occasions. However as several authors have pointed out, and as Elga (2012) acknowledges, once we call on the resources of game theory, we find that several of these rules do entail that a rational agent in Elga's scenario (in which the agent knows that (s)he will be offered both bets) will not reject both bets. See S. Bradley and Steele (2014), Chandler (2014), and Sahlin and Weirich (2014); see Mahtani (2018) for a response.

A further way of responding to Elga's challenge is to argue that when faced with a series of choices, a rational agent will make a plan and stick to it—and where an agent has an imprecise credence function, that plan will be endorsed as maximising expected utility by at least one of the agent's avatars. For further discussion of this sort of view, see Bratman (2012), Gauthier (1986), and McClenen (1990).

Finally, there are authors who reject the assumption that an agent in an Elga-style scenario who rejects both bets is thereby irrational. For example, Moss (2015) constructs an account of what it is for an agent with imprecise credences to "change his or her mind", and argues that it is permissible in at least some Elga-style scenarios for an agent to reject Bet A while identifying with one of her avatars, and then change her mind and reject Bet B, identifying with a different avatar. Others such as S. Bradley and Steele (2014) also maintain that a rational agent in an Elga-style scenario may reject both bets.

Thus there are a range of interesting ways that the imprecise probabilist might respond to the sort of sequential decision problem that Elga

has raised, and the debate over which rule of rationality the imprecise probabilist should endorse is still ongoing.

4 SUMMARY

I began with a natural motivation for accepting imprecise probabilism. I then outlined the most widely discussed account of imprecise probabilities, and considered how the account should be interpreted. I then turned to two categories of objections to the account: objections concerning learning, and objections concerning decision making. Within learning, I discussed two different objections: firstly the problem of belief inertia, and secondly the problem of dilation. Within decision making, I focused on the problems that the imprecise probabilist faces in situations of sequential choice. There has been recent, lively debate about these objections, and while various responses have been put forward by the imprecise probabilists, we are currently far from a consensus.

REFERENCES

- Arrow, K. J. & Hurwicz, L. (1972). An optimality criterion for decision making under ignorance. In C. F. Carter & J. L. Ford (Eds.), *Uncertainty and expectations in economics: Essays in honour of g.l.s. shackle*. Oxford: Basil Blackwell.
- Bradley, R. (2009). Revising incomplete attitudes. *Synthese*, 171(2), 235–256.
- Bradley, R. (2017). *Decision theory with a human face*. Cambridge University Press.
- Bradley, S. & Steele, K. (2014). Should subjective probabilities be sharp? *Episteme*, 11(3), 277–289.
- Bratman, M. (2012). Time, rationality and self-governance. *Philosophical Issues*, 22(1), 73–88.
- Brier, G. W. (1950). Verification of forecasts expressed in terms of probability. *Monthly weather review*, 78(1), 1–3.
- Chandler, J. (2014). Subjective probabilities need not be sharp. *Erkenntnis*, 79(6), 1273–1286.
- Dempster, A. (1967). Upper and lower probabilities induced by a multivalued mapping. *Annals of Mathematical Statistics*, 38(2), 325–39.
- Dempster, A. (1968). A generalization of bayesian inference. *Journal of the Royal Statistical Society*, 30(2), 205–247.
- Dodd, D. (2013). Roger white's argument against imprecise credences. *The British Journal for the Philosophy of Science*, 64(1), 69–77.
- Elga, A. (2010). Subjective probabilities should be sharp. *Philosophers Imprint*, 10(5), 1–11.

- Elga, A. (2012). Errata for subjective probabilities should be sharp.
- Ellsberg, D. (1961). Risk, ambiguity, and the savage axioms. *Quarterly Journal of Economics*, 75(4), 643–669.
- Gärdenfors, P. & Sahlin, N. E. (1982). Unreliable probabilities, risk taking and decision making. *Synthese*, 53(3), 361–386.
- Gauthier, D. (1986). *Morals by agreement*. Oxford: Clarendon Press.
- Gilboa, I. & Schmeidler, D. (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2), 141–153.
- Greaves, H. & Wallace, D. (2006). Justifying conditionalization: Conditionalization maximizes expected epistemic utility. *Mind*, 115(459), 607–632.
- Halpern, J. Y. (2003). *Reasoning about uncertainty*. Cambridge: MIT Press.
- Hammond, P. (1988). Orderly decision theory. *Economics and Philosophy*, 4(2), 292–297.
- Jeffrey, R. (1965). *The logic of decision*. Chicago: University of Chicago Press.
- Jeffrey, R. (1983). Bayesianism with a human face. In J. Earman (Ed.), *Testing scientific theories* (pp. 133–156). University of Minnesota Press.
- Joyce, J. M. (1998). A nonpragmatic vindication of probabilism. *Philosophy of Science*, 65(4), 575–603.
- Joyce, J. M. (2005). How probabilities reflect evidence. *Philosophical Perspectives*, 19, 153–178.
- Joyce, J. M. (2010). A defense of imprecise credences in inference and decision making. *Philosophical Perspectives*, 24(1), 281–323.
- Kaplan, M. (1996). *Decision theory as philosophy*. Cambridge University Press.
- Keynes, J. M. (1921). *Treatise on probability*. London: Macmillan.
- Kyburg, H. (1983). *Epistemology and inference*. Minneapolis: University of Minnesota Press.
- Levi, I. (1974). On indeterminate probabilities. *Journal of Philosophy*, 71(13), 391–418.
- Lewis, D. (1980). A subjectivist's guide to objective chance. In R. C. Jeffrey (Ed.), *Studies in inductive logic and probability* (pp. 83–132). University of California Press.
- Mahtani, A. (2018). Imprecise probabilities and unstable betting behaviour. *Noûs*, 52(1), 69–87.
- Mayo-Wilson, C. & Wheeler, G. (2016). Scoring imprecise credences: A mildly immodest proposal. *Philosophy and Phenomenological Research*, 92(1), 55–78.
- McClenen, F. (1990). *Rationality and dynamic choice: Foundational explorations*. Cambridge University Press.
- Moss, S. (2015). Credal dilemmas. *Noûs*, 49(4), 665–683.
- Pederson, A. & Wheeler, G. (2014). Demystifying dilation. *Erkenntnis*, 79(6), 1305–1342.

- Pettigrew, R. (2013). A new epistemic utility argument for the principal principle. *Episteme*, 10(1), 19–35.
- Rinard, S. (2013). Against radical credal imprecision. *Thought*, 2(1), 157–165.
- Rinard, S. (2015). A decision theory for imprecise probabilities. *Philosophers' Imprint*, 15(7), 1–16.
- Sahlin, N. E. & Weirich, P. (2014). Unsharp sharpness. *Theoria*, 80(1), 100–103.
- Satia, J. & Lave, R. (1973). Markovian decision processes with uncertain transition. *Operations Research*, 21(3), 728–740.
- Savage, L. (1954). *The foundations of statistics*. Wiley.
- Schoenfield, M. (2017). The accuracy and rationality of imprecise credences. *Noûs*, 51(4), 667–685.
- Seidenfeld, T., Schervish, M. J., & Kadane, J. B. (2012). Forecasting with imprecise probabilities. *International Journal of Approximate Reasoning*, 53(8), 1248–1261.
- Seidenfeld, T. & Wasserman, L. (1993). Dilation for sets of probabilities. *The Annals of Statistics*, 21(3), 1139–1154.
- Shafer, G. (1976). *A mathematical theory of evidence*. Princeton, NJ: Princeton University Press.
- Sturgeon, S. (2008). Reason and the grain of belief. *Noûs*, 42(1), 139–165.
- Sturgeon, S. (2010). Confidence and coarse-grained attitudes. In T. S. Gendler & J. Hawthorne (Eds.), *Oxford studies in epistemology*. Oxford: OUP.
- Troffaes, M. (2007). Decision making under uncertainty using imprecise probabilities. *International Journal of Approximate Reasoning*, 45, 17–29.
- Vallinder, A. (2018). Imprecise bayesianism and global belief inertia. *The British Journal for the Philosophy of Science*, 69(4), 1205–1230.
- van Fraassen, B. C. (1984). Belief and the will. *Journal of Philosophy*, 81(5), 235–256.
- van Fraassen, B. C. (1990). Figures in a probability landscape. In M. Dunn & K. Segerberg (Eds.), *Truth or consequence* (pp. 345–56). Amsterdam: Kluwer.
- van Fraassen, B. C. (2006). Vague expectation value loss. *Philosophical Studies*, 127(3), 483–491.
- Voorhoeve, A., Binmore, K., Stefansson, A., & Stewart, L. (2016). Ambiguity attitudes, framing, and consistency. *Theory and Decision*, 81(3), 313–337.
- Walley, P. (1991). *Statistical reasoning with imprecise probabilities*. Chapman & Hall.
- Weatherson, B. (1998). Decision making with imprecise probabilities. Retrieved April 5, 2019, from <http://brian.weatherson.org/vdt.pdf>

- Weatherson, B. (2003). From classical to intuitionistic probability. *Notre Dame Journal of Formal Logic*, 44(2), 111–123.
- White, R. (2009). Evidential symmetry and mushy credence. In T. S. Gendler & J. Hawthorne (Eds.), *Oxford studies in epistemology* (pp. 161–186). Oxford University Press.
- Williams, J. R. G. (2014). Decision-making under indeterminacy. *Philosophers' Imprint*, 14(4), 1–34.
- Yager, R. & Liu, L. (Eds.). (2008). *Classic works of the Dempster-Shafer theory of belief functions*. Studies in Fuzziness and Soft Computing. Springer.