http://eprints.lse.ac.uk

The dispositional account of credence

Anna Mahtani[1]

Abstract In this paper I offer an alternative—the 'dispositional account'—to the standard account of imprecise probabilism. Whereas for the imprecise probabilist an agent's credal state is modelled by a set of credence functions, on the dispositional account an agent's credal state is modelled by a set of sets of credence functions. On the face of it, the dispositional account looks less elegant than the standard account—so why should we be interested? I argue that the dispositional account is actually simpler, because the dispositional choice behaviour that fixes an agent's credal state is faithfully depicted in the model of that agent's credal state. I explore some of the implications of the account, including a surprising implication for the debate over dilation.

Keywords Imprecise probabilism · Credence · Dilation · Rationality

1 Introducing the dispositional account

In this paper, I offer an alternative (the 'dispositional account') to the standard account of imprecise probabilism, and I argue that the dispositional account does as well or better than standard imprecise probabilism on all counts. In this introductory section, I sketch the orthodox Bayesian view and then introduce both the standard imprecise probabilism account and my own dispositional account.

On the orthodox Bayesian view, a rational agent's credal state can be represented by a credence function that maps every proposition (s)he can entertain onto some

✉ Anna Mahtani
a.mahtani@lse.ac.uk

1 Department of Philosophy, Logic and Scientific Method, London School of Economics and Political Science, Houghton Street, London WC2A 2AE, UK

Springer

number between 0 and 1.[1] For many propositions, this seems reasonable: plausibly you have a credence of exactly 0.5 that (HEADS) a coin that you know to be fair will land heads on its next toss. However, there are also many propositions for which the orthodox Bayesian view doesn't seem plausible. For example, take the following proposition:

SARDINES: Anna Mahtani's neighbour has at least one tin of sardines in her cupboard.

What precisely is your credence in SARDINES?[2] I'm guessing that no particular number springs to mind. If pushed you might produce a number, but the number produced would be arbitrary: you wouldn't *know* that it is your precise credence in SARDINES. Explaining your ignorance here is one challenge for the orthodox Bayesian. Another is to answer this question: what could make any particular precise value your credence in SARDINES? For example, what could make it the case that your credence in SARDINES is precisely 0.356, say, rather than 0.357?

In response to these and other[3] problems with orthodox Bayesianism, many have embraced imprecise probabilism. There are numerous versions of imprecise probabilism that can be found in the literature, including the theory of lower previsions, partial probability orderings, and sets of desirable gambles (Walley 1991; Augustin et al. 2014). Here for simplicity I focus on one framework from the literature, on which an agent's credal state is represented not by one credence function (as it is for the orthodox Bayesian) but instead by a set of credence functions (Levi 1974; Jeffrey 1983; Joyce 2005). For example, your credal state might be represented by a set of credence functions which between them assign every number strictly between 0.1 and 0.8 to SARDINES. We can then say that your credence in SARDINES is the range (0.1, 0.8).[4]

Having set out (in rough outline) the orthodox Bayesian view and imprecise probabilism, I turn now to the new alternative that I am proposing—the dispositional account. To motivate the account, I start by considering the following passage from De Finetti:

> One can … give a direct, quantitative, numerical definition of the degree of probability attributed by a given individual to a given event … It is a question simply of making mathematically precise the trivial and obvious idea that the degree of probability attributed by an individual to a given event is revealed by the conditions under which he would be disposed to bet on that event.
>
> Let us suppose that an individual is obliged to evaluate the rate p at which he would be ready to exchange the possession of an arbitrary sum S (positive or

---

[1] One introduction to this orthodox Bayesian view can be found in (Skyrms 2000, pp. 137–144).

[2] This—like Adam Elga's 'toothpaste-in-the-bag example', is a case where evidence for the proposition is 'sparse and unspecific' (Elga 2010, pp. 1–2).

[3] Another reason to move from orthodox Bayesianism to imprecise probabilism comes from the thought that adopting an imprecise credence can be the only rational response to the evidence available (Joyce 2005, p. 171). Some further related reasons to reject orthodox Bayesianism can be found in (Bradley 2009).

[4] Note though that the model of your credal state (consisting of a set of credence functions) will contain more information than just the range assigned to each individual proposition (Joyce 2010, pp. 285–287).

negative) dependent on the occurrence of a given event E, for the possession of the sum pS; we will say by definition that this number p is the measure of the degree of probability attributed by the individual considered to the event E…

(De Finetti 1964, pp. 101–102)

Let us consider how this might work for the proposition (HEADS) that a coin that you know to be fair will land heads on its next toss. What number would you produce if I were to follow De Finetti's method to elicit your rate p for this proposition? The trick to De Finetti's method is to require you to produce your rate p before the stake is settled, and so you need to select the number you produce with care to avoid ending up committed to a bet that is a bad deal by your own lights.[5] We say that the number you produce (presumably 0.5 in this case) is your credence in HEADS.

Of course, you can have a credence in a proposition even if you have not actually been subjected to De Finetti's test. The idea is that your credence is equal to the number that you *would* produce *were* you to be subject to such a test. We can see this subjunctive mood in De Finetti's own words—'the conditions under which he would be disposed to bet…' (De Finetti 1964, p. 101). How should we interpret this subjective claim? Let us combine De Finetti's account with a popular philosophical account of counterfactuals, broadly based on David Lewis' analysis (Lewis 1973).[6] On this account, to ask what *would* have happened had some event occurred, is to ask what happens in those closest possible worlds (or 'counterfactual scenarios') where the event *does* occur. Thus if you have not actually been subjected to De Finetti's test over HEADS, then instead of looking for some number that you produce in the actual world we should instead look for the number or numbers that you produce in the closest world or worlds where you are subjected to De Finetti's test over HEADS. Plausibly there are a range of equally close relevant worlds, but provided that we focus on those close worlds where you have the same relevant evidence and rationality as you have in the actual world, you would presumably produce the same number (0.5) in each. Thus we can define your credence in HEADS as *the* number that you would produce were you subjected to De Finetti's test.

We get a different picture when we switch our attention to SARDINES. Consider the range of closest worlds in which you are subjected to De Finetti's test over the proposition SARDINES: we would expect the number that you produce to vary across these close worlds. The reason for this is that the number that you produce in each of these worlds is to some extent arbitrary: you may be determined not to state a number below 0.1 or above 0.8, but you have no particular reason to produce the number 0.356 rather than 0.357. Thus there will be some worlds where you produce the number 0.356, and others where—with no variation in your evidence, nor in the

---

[5] For a helpful explanation of De Finetti's method, see (Gillies 2000, pp. 55–58).

[6] Lewis' account of counterfactuals has faced objections and undergone various refinements—but these issues are largely orthogonal to my argument in this paper.

terms of the bet, and with no failure of rationality—you instead produce the number 0.357. You have to pick a number somehow, so you will decide on a whim, or perhaps even rely on some randomized process to help you choose, and in different close possible worlds your whim might pull you in a different direction, or the random process might have a different result. I have argued elsewhere (Mahtani 2018) that this instability in your betting behaviour is characteristic of those cases which motivate a move from orthodox Bayesianism to imprecise probabilism.

Thus there is no such thing as *the* number that you would produce were you subjected to De Finetti's test over SARDINES. So if we define your credence in SARDINES as the number that you would produce in these circumstances, then we arrive at the conclusion that there is no such thing as your credence in SARDINES. But there is a natural way to extend De Finetti's definition to avoid this result.

De Finetti's claim, together with our account of counterfactuals, entails that an agent's credence in a proposition is fixed by the choice that he or she makes in those closest counterfactual scenarios where (s)he is subjected to De Finetti's test. For some propositions like HEADS, it is plausible to assume that the agent makes the same choice across these counterfactual scenarios, and so this single choice can define the agent's credence in that proposition. For other propositions like SARDINES, this assumption is not plausible: an agent does not make a single choice across all of these scenarios, but makes a range of different choices. We can say then that an agent's credal state is defined by the choices that he or she makes across these relevant counterfactual scenarios. This is the rough idea behind the dispositional account of credences—which I will go on to formalize in the next section.[7]

Before moving onto the formalization, I pause to clarify which counterfactual scenarios count as 'relevant' on the account. For this paper I count as relevant just those close counterfactual scenarios in which an agent is subjected to De Finetti's test, and in which the agent has the same evidence and preferences as (s)he actually has and suffers no loss of rationality. On my dispositional account, an agent's credal state is defined by his or her choice behaviour in these relevant scenarios. But some might argue that an agent's credal state is not fixed merely by his or her dispositional *betting* behaviour: for consider, an agent might have a moral objection to betting, and so refuse to produce any number at all in response to De Finetti's test or produce a number at random. In such a case (the objector might say), the agent's credal state is not reflected in his or her (overt) betting behaviour,[8] and cannot be defined by it.[9]

---

[7] On my view, a key characteristic of cases where intuitively credences are imprecise is that our choice behaviour across close possible worlds is unstable. Offered the very same bet, on the very same terms, we may choose differently in different close possible worlds—and with no good reason to explain the difference. This does not make us irrational. In contrast to my view, many other views in the literature characterise cases where credences are imprecise as those where there is a difference between the price at which one would sell a bet and the price at which one would buy the same bet. As my account is not concerned with this difference, it is only tangentially related to work by Walley and others on sets of desirable gambles (1991).

[8] I say 'overt' betting, because you might agree with Ramsey that 'all our lives we are in a sense betting' (Ramsey 1931, p. 183).

[9] These and other objections to accounts that define an agent's credal state in terms of his or her actual or dispositional betting behaviour are raised in (Eriksson and Hajek 2007).

The natural (and I think right) solution to this problem is to extend the range of relevant counterfactual scenarios to include scenarios where the agent makes choices that do not involve (overt) betting at all. Thus for example an agent's epistemic attitude towards the proposition that his train is about to leave is not just defined by his or her choice behaviour in scenarios where (s)he is asked to bet on that proposition, but also by his or her choice behaviour in scenarios where (s)he is faced with other choices relating to that proposition—such as the choice between walking and running towards the platform. If we were to extend the account in this way, then which choice scenarios exactly should we include? A rough answer is that we should include all choice scenarios in which the agent's mental state is relevantly similar to his or her mental state in the actual world: for example, worlds in which (s)he has the same relevant evidence, preferences, and powers of rationality as in the actual world. In developing a more precise answer, we should be guided by the goal of fitting broadly into the more general theory of functionalism. According to functionalism, a mental state is defined by its functional role—by its role within a complex network which includes both sensory inputs and behavioural outputs. The counterfactual choice scenarios that feature in the dispositional account can be seen as specifications of the agent's output (his or her choice behaviour) in response to certain inputs (choice problems). Of course, for the functionalist many other sorts of inputs and outputs will feature. Furthermore, the functionalist account is holistic—and so any definition of a single mental state may need to make reference to every other mental state. An extended version of the dispositional account would undoubtedly still be a simplification, but it should be broadly compatible with the functionalist project.

Thus there are ways that the dispositional account could be developed, but in this paper I have a narrower focus. The goal is just to extend the orthodox Bayesian account (which traditionally works with a concept of credal state that is closely tied to betting behaviour) in a way that is better—more informative, and more faithful to what is being modelled—than standard imprecise probabilism. Thus I limit the relevant counterfactual scenarios to those in which agents are subjected to De Finetti tests, and I set aside the worry that an agent's credal state cannot be defined in terms of this sort of betting behaviour. The possibility of developing the account further is left open but the aim of this paper is to produce workable models that improve on those produced by standard imprecise probabilism.

Having outlined the basic thought behind the account, I now offer a formalization.

2 A formalization

The rough idea is that an agent's credal state is defined by the choices that the agent makes in a range of relevant counterfactual scenarios where (s)he is subjected to one or more De Finetti style tests. I start by explaining how we can represent the choice or choices that an agent makes in a single such scenario: the idea is that we represent the agent's choice behaviour at a single scenario with the set of complete and precise credence functions that fit that choice behaviour—i.e. that assign the

right values to the relevant propositions. So for example, if we take a scenario in which an agent is subjected to a De Finetti test over the proposition HEADS and produces the number 0.5, then we can represent the agent's choice behaviour at this scenario by the set containing all and only those credence functions that assign 0.5 to HEADS.

Why does an agent's choice behaviour at a scenario get represented by a *set* of complete and precise credence functions, rather than just by a single credence function? The reason is that in a single choice scenario an agent may be subjected to De Finetti tests over one or more propositions, but at a single choice scenario an agent need not subjected to a test over *every* proposition. For example, in a given choice scenario an agent might only be subjected to a De Finetti test over HEADS and produce the number 0.5. There are many credence functions that assign the number 0.5 to HEADS, and so the agent's choice behaviour at this scenario is represented by the set of all of these credence functions.[10]

Thus an agent's choice behaviour at a particular relevant scenario is represented by a set of complete and precise credence functions. And the idea behind the dispositional account is that an agent's credal state is defined by his or her choice behaviour at the range of relevant scenarios—that is, the range of close scenarios in which (s)he is subjected to one or more De Finetti tests. This suggests a formal way to represent an agent's credal state. For each relevant scenario, we take the set of credence functions that represent the agent's choice behaviour at that scenario, and then we gather all these sets of credence functions into a set—which we can call a 'mega-set'. The mega-set represents the agent's dispositional betting behaviour, and so this is a model of the agent's credal state. Thus an agent's credal state is represented by a mega-set containing sets of credence functions.

At this point there is a choice to make: should we model an agent's credal state with a mega-set containing just sets of *probability* functions? Or should we allow in other non-probabilistic credence functions too? On my view, we should allow in non-probabilistic credence functions. This is perhaps the most controversial part of the dispositional account, so I pause here to discuss this point in more depth. Why wouldn't we construct an account free from this controversial claim? On such an account, the agent's behaviour at each choice scenario would be represented by a set of probability functions, and thus an agent's credal state would be represented by a mega-set of sets of probability functions, and no non-probabilistic credence functions would feature. The drawback is that this alternative account produces models that are less informative than the version of the dispositional account that I recommend. I demonstrate that this holds for models of irrational agents in Sect. 5, and for models of rational agents in Sect. 8. So on my preferred version of the

---

[10] I could have set up the account differently, so that there is just one credence function corresponding to each scenario, but then that credence function would be only partial. A partial credence function can be represented by the set of precise credence functions that 'extend' it, so this may be another route to the same sort of account. But to end up at the same account, with the same advantages, we would need to think of the partial credence function as representable by a set of credence functions that might include non-probabilistic credence functions. This contrasts with Richard Bradley's account of partial credence functions, on which a partial credence function can be represented by the set of probability functions that extends it (Bradley 2017, pp. 371–375).

dispositional account we model the credal state of an agent—even a rational agent—using a mega-set of sets of credence functions, some of which may not be probability functions.

An account that introduces credence functions that are not probability functions may seem very unconventional. However it is undeniable that there are functions that are not probability functions, but assign values to the very same propositions as some probability functions do. To be more precise, for any given algebra that is the domain for a probability function, there will be some other non-probabilistic function (in fact infinitely many such functions) that have the same domain.[11] For example, if we have a probability function that (amongst other assignments) assigns 0.4 to P, 0.3 to Q, 0 to (P∧Q), and 0.7 to (P∨Q), there will be a non-probabilistic function that assigns the same numbers to P, Q and (P∧Q) but, say, 0.8 to (P∨Q). Thus there are certainly functions that can play the role that I have in mind. Whether they can legitimately be called 'credence functions', or whether they should be called something else ('quasi-credence functions', perhaps) is open to debate— though non-probabilistic credence functions seem essential if we want to model the epistemic states of agents who are less than perfectly rational. At any rate, in this paper I will persist in calling them 'credence functions' on the understanding that nothing turns on this name.[12,13]

Having set the dispositional account out more formally, I turn now to explain why you should be interested in this account when the standard imprecise probabilist account is already available.

3 How could this be an improvement on imprecise probabilism?

On the face of it, standard imprecise probabilism is simpler and more elegant than the new dispositional account that I am proposing. After all, on the standard account an agent's credal state is represented by a set of credence functions, whereas on the dispositional account an agent's credal state is represented by a mega-set of sets of credence functions, which certainly sounds more complicated. However, I argue that the dispositional account is in fact simpler.

---

[11] There are further questions to be raised about what sorts of things the algebra and state-space should be. For example, should there be a credence function that assigns one value to some proposition P, and another value to ¬¬P? This will depend on whether 'P' and '¬¬P' express the same element in the algebra, and this in turn may depend on what sorts of things are the underlying states. This is an interesting and important question, but not one that I delve into here.

[12] Thanks to Nicholas Baigent for challenging me on this point.

[13] Wlodek Rabinowicz argues that where an agent's credence function is not probabilistic, we may not be able to use the claim that the agent is maximising expected utility to *explain* why the agent's response to a test like De Finetti's will equal his or her credence. He writes: 'If an agent violates some of the basic assumptions of the expected-utility theory (such as the standard probability axioms), then we cannot explain her behaviour by an appeal to expected-utility considerations.' (Rabinowicz 2014, p. 361). As justifying the agent's choice behaviour is not the focus of this paper, I follow Rabinowicz's lead and set this point aside here.

To see that imprecise probabilism is not as simple as it appears, we can interrogate the account beginning with this question: according to imprecise probabilism, an agent's credal state is modelled by a set of precise credence functions, but what in reality do these precise credence functions each represent? In response, the imprecise probabilist might reasonably retort that I am misunderstanding the account: the credence functions in the set do not individually represent anything in reality. The idea is rather that they jointly represent the credal state of the agent. For example, a range of credence functions between them assigning to SARDINES every number between 0.1 and 0.8 are members of the set that represents your credal state because your credence in SARDINES is the range (0.1, 0.8). We can then turn our attention to this question: what fixes facts about an agent's credal state? What makes it the case, for example, that your credence in SARDINES is the range (0.1, 0.8)? In response, the typical imprecise probabilist will give a broadly functionalist response: your credence in SARDINES is the range (0.1, 0.8) because this reflects the choices that you would make in various counterfactual scenarios. Thus we can see that imprecise probabilism has a three-part structure: first we have your choice behaviour in various counterfactual scenarios; second we have your credal state, which is somehow fixed by this choice behaviour; and third we have a model that represents your credal state—a model that consists in a set of precise credence functions. We might attempt to complete the circle by matching up one-to-one each precise credence function in the set with your choice behaviour at some counterfactual scenario—but this is not part of the account and it is not obvious how this should be done.[14] Thus we are left with an account with a three-part structure.

The structure of the dispositional account is simpler. On this view—as for imprecise probabilism—your credal state is fixed by your choice behaviour (specifically, your betting behaviour) in a range of counterfactual scenarios. But—unlike imprecise probabilism—the dispositional account directly represents your choice behaviour in each counterfactual scenario with a set of credence functions. This gives us a model of your credal state that consists in a mega-set of sets of credence functions. If we now interrogate the dispositional account by asking the question 'where do all these sets of credence functions come from?', then we have a clear answer: each set of credence functions in the mega-set represents your choice behaviour at some counterfactual scenario—and we can spell out this representation relation precisely.

Despite first appearances, then, the dispositional account is simpler than imprecise probabilism: the choice behaviour that defines an agent's credal state is

---

[14] On many of the decision rules that have been proposed by imprecise probabilists, it follows that every choice made by an agent will—if the agent is rational—maximise expected utility relative to at least one credence function in the set that represents the agent's credal state. Presumably then we could define a relation that maps each choice made by the agent to some credence function in the set that represents the agent's credal state. But often there will be a range of credence functions in the set relative to which the choice maximises expected utility, and so a choice could be mapped equally appropriately to many different credence functions, and to choose between them would be arbitrary. We might instead then map each choice onto a set of credence functions—and this takes us back to the dispositional account that I am proposing.

directly represented in the model that represents that agent's credal state. Later in the paper (in Sect. 8) I argue that the dispositional account has another advantage over imprecise probabilism: it produces models that are more informative. First however I turn to develop the dispositional account.

4 Developing the account (i)—rationality

What is required for an agent to count as synchronically rational—as far as his or her credal state is concerned?[15] For the imprecise probabilist, the answer is this: the set of credence functions that represent the agent's credal state must contain only probability functions (Joyce 2010, p. 287). But what answer should we give for the dispositional account?

To put the answer intuitively first, an agent is rational provided that his or her choice behaviour is coherent at every relevant scenario. And choice behaviour at a scenario is coherent provided that there is a probability function that fits that behaviour.[16] Thus (to put the answer more formally) we arrive at this view: an agent is rational with respect to his or her credal state if and only if within every set (in the mega-set that represents the agent's credal state) there is at least one credence function that is a probability function.

To illustrate we can consider a toy example on which there are just three relevant scenarios: S_1, where you are subjected to De Finetti's test over SARDINES, and produce the number 0.4; S_2, where you are subjected to De Finetti's test over SARDINES, and produce the number 0.6; and S_3, where you are subjected to De Finetti's test over HEADS, and produce the number 0.5. We are imagining that in each of these scenarios you are only faced with one de Finetti style test—though of course that need not be the case in general. Within each scenario your choice behaviour is coherent: that is, some probability function fits it. Thus, in every set of credence functions (within the mega-set that represents your credal state) there will be a probability function, and so you meet the rationality requirement. Note that there is no requirement that you are coherent across the different scenarios: you can be rational even though there is no probability function that fits your choice behaviour at both S_1 and S_2. For you to be classed as irrational, your choice behaviour at a *single* scenario must be such that no probability function fits it.

To illustrate, let us add a further scenario S_4, in which you are subjected to De Finetti's test over both SARDINES and ¬SARDINES, and produce the answers 0.4 and 0.5 respectively. No probability function fits your choice behaviour in S_4, and so there is no probability function in the set that corresponds to this scenario. Thus in the mega set that represents your credal state, there will be at least one set of

---

[15] Here I do not explore the question of what diachronic rationality requires on the dispositional account: this is a topic for another paper.

[16] We might require more (or less) of an agent than this. There are many variations on the account that could be considered. For example, we might rule that choice behaviour at a scenario is coherent provided that the behaviour fits with some probability function for which the principal principle (or any other additional principle of rationality) holds.

credence functions that does not contain a probability function, and so you are classed as irrational.[17]

What is meant by a choice 'scenario'? In a possible world an agent may face one or more De Finetti style betting choices. A string of choices will count as made within the same scenario only if the agent has the same relevant evidence, preferences, and powers of rationality throughout. This is a necessary condition for a string of choices to be classified as being made within the same scenario, but it is not a sufficient condition. To see this, suppose that an agent is offered one De Finetti style betting choice on SARDINES and makes his or her choice, but then learns that the bet has been cancelled; if the agent is then offered another similar betting choice over SARDINES, (s)he is free to choose differently without thereby being irrational. Thus we should not class these two choices as part of the same scenario: intuitively, the slate has been wiped clean between the choices. In judging whether a series of choices made by an agent counts as a single scenario, we should be guided by our judgements about what is rationally permissible.[18] Sometimes we require a string of choices to cohere for an agent to count as rational, and sometimes we do not: relevant factors affecting our judgement include the length of time between the choices being made, whether the outcomes of the choices aggregate, and what information the agent has throughout about the choices that (s)he will face. It is hard to give a precise definition of 'scenario' here, for plausibly the term 'scenario' is vague. Instead, I aim just to carve out a set of clear cases where several conditions hold which are jointly sufficient for a string of choices to count as a scenario. In such clear cases, the choices are offered to the agent in quick succession; the agent knows immediately before the choices are offered exactly which choices (s)he will be offered and in what order; the choices are known to have outcomes that matter to the agent and that aggregate, and the agent has the same preferences throughout the series and neither gains nor loses any powers of rationality or relevant evidence. If the agent is rational, then all of his or her choice behaviour across such a scenario will fit with at least one probability function.

This notion of a scenario is something that standard imprecise probabilists need too, and the challenge of spelling it out more precisely is a challenge that we share. For the standard imprecise probabilist needs to give a decision rule (I say more about decision rules in Sect. 6), and that decision rule must determine what is rationally required of an agent who is faced with a string of choices. Various examples in the literature show that this is not such an easy task (Hammond 1988, Elga 2010), and I argue elsewhere (Mahtani 2018) that the response to this problem must involve the concept of a scenario—or a 'sequence' as it is called in that

---

[17] The version of the dispositional account that I recommend can produce informative models even of irrational agents: a model can help us predict the way that an agent will behave even in scenarios such as S_4 where that agent's behaviour is not coherent. An alternative version of the dispositional account (discussed in Sect. 2) produces models that feature only probability functions (rather than also including non-probabilistic credence functions) and the models of irrational agents produced on this version will be less informative—for at any scenario such as S_4 where an agent's choice behaviour is incoherent, the set of probability functions corresponding to that scenario will be empty.

[18] This does not make the dispositional account circular, for the account is not intended to be reductive.

literature (Elga 2010; Weatherson 2008; Williams 2014). Thus the standard imprecise probabilist also owes us a more precise definition of a scenario: both imprecise probabilism and the dispositional account need to make reference to scenarios in stating what rationality requires of an agent.

Having clarified what it is for an agent to count as rational, I turn now to another development of the account.

5 Developing the account (ii)—claims about the agent

At least some versions of imprecise probabilism offer an account of truth-conditions for claims apparently describing an agent's (single) credence function: such a claim is true if and only if the claim is true of every credence function in the set that represents the agent's credal state.[19] Thus for example, the claim that your credence in SARDINES is less than 0.9 is true provided that every credence function in the set that represents your credal state assigns a number lower than 0.9 to SARDINES. To give another example, the claim that your credence in SARDINES is greater than your credence in EXTREME SARDINES (where EXTREME SARDINES is the claim that Anna Mahtani's neighbour has at least 2 tins of sardines in her cupboard) holds provided that every credence function in the set assigns a higher number to SARDINES than it assigns to EXTREME SARDINES. Of course, according to the imprecise probabilist the best model of your credal state is not a single credence function—but we can understand these claims about your credence function as a sort of shorthand: a claim about your credence function is really a quick way of making a claim about your more complex credal state and the whole set of credence functions that can be used to model it.

Can we make a similar move on the dispositional account? Just as for imprecise probabilism, on the dispositional account the best model of an agent's credal state does not consist of a single credence function. Still, we can similarly make claims apparently about an agent's single credence function as a sort of shorthand: we can give a partial description apparently of your (single) credence function, on the understanding that what we are really saying is something about your more complex credal state and the mega set of sets of credence functions that can be used to model it. Let us say that a claim apparently describing your single credence function is true, provided that your choice behaviour at every relevant scenario fits with some credence function for which the description holds; equivalently, provided that each set (within the mega-set that represents your credal state) contains at least one credence function for which the description holds. Intuitively, the requirement is that you always act (that is—at each relevant scenario you act) as if you had a credence function that meets the description.

As an illustration, let's consider a toy example from the last section. In this example, there are just three relevant choice scenarios: S_1, where you are subjected

---

[19] This is not quite precise, for I have not defined what it would be for a claim to be 'true of' some credence function—but the examples should make it clear.

to De Finetti's test over SARDINES and produce the number 0.4; S_2, where you are subjected to De Finetti's test over SARDINES and produce the number 0.6; and S_3 where you are subjected to De Finetti's test over HEADS and produce the number 0.5. Now we can consider this claim which apparently describes your (single) credence function: your credence in SARDINES is less than 0.5. Your choice behaviour in S_1 and S_3 fits with credence functions for which the description holds, but your choice behaviour in S_2 does not. To see this, consider the set of credence functions that represent your choice behaviour in S_2: every one of them assigns the number 0.6 to SARDINES (though they can differ in the number that they assign to other propositions[20]). Thus in this set there is no credence function that meets the description, and so the claim that your credence in SARDINES is less than 0.5 does not hold. Now let us consider instead this claim about your credal state: your credence in SARDINES is less than 0.9. We can see that your choice behaviour in all three scenarios fits with credence functions for which the analogous claim holds. The set of credence functions that represent S_1 all assign 0.4 to SARDINES—and so all of them meet the description (and so of course at least one of them does). Similarly the set of credence functions corresponding to S_2 all assign 0.6 to SARDINES—and so all meet the description. Finally the set of credence functions corresponding to S_3 assign a whole range of different numbers to SARDINES— including numbers below 0.9. This is because in S_3 you are asked to make a choice that just concerns the proposition HEADS, and any credence function that assigns 0.5 to HEADS fits with your behaviour at this scenario—regardless of the number assigned to SARDINES. Thus in the set of credence functions that represent your choice behaviour at S_3, there will be at least one credence function that meets the description. Thus we can say (bearing in mind of course that this is just a sort of shorthand) that your credence in SARDINES is less than 0.9.

Let us now consider a slightly more complex claim, involving two propositions: the claim that your credence in SARDINES is higher than your credence in EXTREME SARDINES. Let us use the following toy example, again with just three scenarios. In one (S_5), you are subjected to De Finetti's test over SARDINES and over EXTREME SARDINES and produce the numbers 0.4 and 0.3 respectively. In another (S_6), you are again subjected to De Finetti's test over SARDINES and EXTREME SARDINES and produce the numbers 0.6 and 0.5 respectively. In the other scenario (S_7) you are subjected to De Finetti's test over HEADS and produce the number 0.5. The claim under consideration (that your credence in SARDINES is higher than your credence in EXTREME SARDINES) holds, because your choice behaviour in every scenario fits with some credence function that meets the description. For scenarios S_5 and S_6 all the credence functions that fit your behaviour meet the description, and for scenario S_7, some of the credence functions that fit with your choice behaviour meet the description.

Of course, these toy examples are absurdly simple in lots of ways. The number of relevant choice scenarios is far too small: your credal state will depend on the

---

[20] On my preferred version of the dispositional account (on which models can include non-probabilistic credence functions), this even includes some propositions that are logically related to SARDINES.

choices that you would make in many more scenarios than that.[21] These toy examples are just designed to show how, on the dispositional account, we can make a claim apparently about an agent's (single) credence function as a shorthand for a claim about his or her more complex credal state.

I turn now to outline the decision theory on the account.

6 Developing the account (iii)—decision theory

Imprecise probabilists have put forward many different decision theories, and here I describe just one sample theory, known as 'E-admissibility' (Levi 1974) or 'caprice' (Weatherson 2008). According to this theory, an agent faced with a choice problem is rationally required to maximise expected utility relative to at least one credence function in the set that represents his or her credal state.[22] Thus for example, suppose that you are subjected to De Finetti's test over the proposition SARDINES, and let us assume that you value only money, and value that linearly. Then you are permitted to produce the number 0.35, provided that this is the value assigned to SARDINES by some credence function in the set that represents your credal state, because producing the number 0.35 maximises expected utility relative to that credence function. Similarly, you are permitted to produce the number 0.34, provided that some other credence function in the set assigns the number 0.34 to SARDINES—and so on.

What about a decision theory for the dispositional account? Here is one possible answer. Within any relevant scenario, an agent makes choices that fit with *all* the credence functions within the set that represents his or her choice behaviour at that scenario. We could require this to hold more generally: within any choice scenario (whether it concerns overt betting or not), an agent should make choices that fit with *all* the credence functions within *some* set in the mega set that represents his her credal state.

---

[21] We could assume that for any combination of propositions that you can entertain, there is some relevant scenario in which you are subjected to a De Finetti style test over every member of that combination of propositions. Without this assumption, we get some weird results. Take this toy example as an illustration. There are two choice scenarios: in one (S_8) you are subjected to a De Finetti test just over proposition P and return the number 0.4; in the other (S_9) you are subjected to a De Finetti test just over proposition (P∨Q) and return the number 0.3. You count as rational, because your choice behaviour at every scenario fits with some probability function, and yet the following claim (apparently) about your single credence function comes out as true: you have a credence of 0.4 in P, and a credence of 0.3 in P∨Q. We can avoid this problem by making the assumption that there must be some relevant scenario in which your credence in both P and (P∨Q) is tested, for here you cannot produce the number 0.4 for P and 0.3 for (P∨Q) and have your choice behaviour fit with some probability function, and so we avoid the the weird result. But we do not *have* to make this assumption: provided that we keep in mind that claims about an agent's single credence function are just a sort of shorthand and need to be unpacked, then we can see that the result is not really weird after all.

[22] Many imprecise probabilists require that an agent's choice behaviour across a *sequence* of choices must maximise expected utility relative to at least one credence function in the set representing their credal state (see Sect. 4).

But is this decision theory normative or descriptive? The make-up of the mega set depends simply on the agent's dispositional choices, and then which choices the decision theory permits the agent to make depends just on the make-up of the mega set. In fact on any broadly functionalist (or behaviourist) account, there will be circularity of this sort. For example, for the typical imprecise probabilist it is facts about actual and dispositional choice behaviour that determine which credence functions belong in the set that represents the agent's credal state, and then the choices that the agent is permitted to make depend on which functions are in this set. Thus both the dispositional account and the typical imprecise probabilist's account are committed to this sort of circularity. It is just that the circularity is very stark on the dispositional account because this account cuts out an intermediate step and effectively represents an agent's credal state directly by the agent's actual and dispositional choice behaviour. On this account, it is the rationality requirement (as described in Sect. 5 above) that is normative. This rationality requirement entails that in any given choice scenario, a rational agent will act in a way that is endorsed by some probability function.

I have now developed the account in three ways: I have stated the rationality requirement, explained how we can read off facts about an agent from a model of that agent, and given a decision theory for the account. I turn now to draw out a surprising implication, which is that the problem of dilation vanishes on the dispositional account.

7 Dilation

I introduce dilation using an example from White (2010).[23] Take some claim P that you have no evidence whatsoever either for or against. Let us assume that on the imprecise probabilist's account your credence in P is the range (0,1): in other words, the credence functions that represent your credal state assign between them every number strictly between 0 and 1 to P. Now suppose that I know whether P is true, and I take a coin that we both know to be fair and paint the heads side over so that the heads symbol is not visible. Without you seeing, I write 'P' on this heads side if and only if P is true, and '¬P' on the heads side if and only if ¬P is true. I similarly paint over the tails side of the coin and write on this side whichever claim (out of 'P' and '¬P') is false. You know that I have done this—that is, you know that the true claim is on the heads side, and the false claim is on the tails side—but you have not learnt anything about the likelihood of P from this process so your credence in P is still (0,1). You know that I am now about to toss the coin, and your credence at t_0 just before it is tossed that (HEADS) it will land heads is 0.5. I toss the coin before your eyes, and at t_1 you see it land P side up. What then at t_1 is your credence in P and your credence in HEADS?[24]

---

[23] Earlier discussions of dilation can be found in (Walley 1991; Seidenfeld and Wasserman 1993).

[24] Arthur Pedersen and Gregory Wheeler have argued for a distinction between proper and improper dilation, and further that White's example is a case of improper dilation (Pederson and Wheeler, 2014). I focus on White's example in this paper, but the points could equally well be made by taking a case that

At t_1 you have learnt that the coin has landed P side up. Thus you know that if P is true, then HEADS is also true—for (you can reason) 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 side up. Furthermore, you know that if HEADS is true, then P is also true—for (you can reason) if it has landed heads side up 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 can be certain that (P iff HEADS). It seems then that 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 (0.5) and your credence in P (0,1) were different, 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.[25,26] It seems though, as Joyce writes, that this phenomenon 'is entirely unavoidable on any view that allows imprecise probabilities' (Joyce 2010, p. 299). As I show, however, the phenomenon is not inevitable on the dispositional account.

To show this, I start by zooming in on a move in the argument above. This is the move from the claim that at t_1 you are certain that (P iff HEADS) to the claim that you must therefore have the same credence at t_1 in HEADS as you have in P. This follows on the imprecise probabilist's account. To see why, suppose that at t_1 you have a credence of 1 that (P iff HEADS), and a credence of 0.5 in HEADS. It follows that every credence function in the set that represents your credal state at t_1 assigns 1 to (P iff HEADS), and 0.5 to HEADS. Thus if every credence function in the set is a probability function (which must be the case if you are rational), then every credence function must assign 0.5 to P, in which case your credence in P is 0.5. More generally, if every credence function assigns a value of 1 to (P iff HEADS), then every credence function must assign the same value to HEADS as it assigns to P, and so inevitably your credence in HEADS will be the same as your credence in P. This step of the argument goes through for the imprecise probabilist, then, but the analogous step does not go through on the dispositional account. On the dispositional account, you can be rational, have a credence of 1 in (P iff HEADS), a credence of 0.5 in HEADS, and yet not have a credence of 0.5 in P. To show how this can be done, I first give a toy example in which this holds, and then I give an intuitive gloss.

Here is the toy example. Let's suppose that the relevant counterfactual choice scenarios all involve you being subjected to De Finetti style tests over some

---

combination of these three claims: HEADS, (P iff HEADS), and P. In any scenario where you are subjected to a De Finetti test over HEADS you produce the number 0.5; in any scenario where you are subjected to a test over (P iff HEADS) you produce the number 1. But your response to tests over P is slightly more complicated. In scenarios where you are subjected to a test over P and also tests over HEADS and/or (P iff HEADS), you produce the number 0.5 in response to the test over P. But across scenarios where you are subjected to a test *just* over P, the number you produce varies strictly between 0 and 1: that is, for every number strictly between 0 and 1, there is some scenario in which you are subjected to a test just over P and produce that number. As described in Sect. 5, we can then make various claims, apparently about your single credence function, as a sort of shorthand. In this case we can say that your credence in HEADS is 0.5, because at every relevant scenario, your choice behaviour fits with a credence function that assigns 0.5 to HEADS. Similarly, we can say that your credence in (P iff HEADS) is 1. But we cannot say that your credence in P is 0.5, because there are relevant scenarios where your choice behaviour does not fit with any credence function that assigns 0.5 to P. For example, take the scenario in which you are subjected to a test over just P, and return the number 0.35. Every credence function that fits with this choice behaviour will assign 0.35 to P, and so there will be no credence function in this set that assigns 0.5 to P. More generally, we cannot claim that your credence in P is any particular number: all we can say is that it is larger than 0, and smaller than 1. Furthermore, you meet the rationality requirement, for your choice behaviour at every individual scenario is endorsed by some probability function. On the dispositional account, then, to put it roughly, your credence in P and in HEADS may be the same at t_1 as they were in t_0: there is no need for your credence in HEADS to dilate, nor for your credence in P to become sharp.

This is just a toy model to show how it is theoretically possible to avoid the problem of dilation on the dispositional account. But it also fits with how we might reasonably expect an agent to behave in White's scenario. At t_1 after seeing the coin land P-side up, the agent is certain of the claim (P iff HEADS) and would bet accordingly. The agent's credence in HEADS is (intuitively) unchanged between t_0 and t_1: seeing the coin land P-side up does not make the agent think that HEADS is any more or less likely than before, so the agent would continue to bet at 50-50 odds on HEADS.[27] But similarly, seeing the coin land P-side up does not give the agent any useful evidence about the likelihood of P, so the agent would continue to exhibit unstable choice behaviour over P: if subjected to De Finetti's test just over P, the agent may respond in a way that is compatible with having any credence strictly between 0 and 1 in P.[28] However, if the agent's credence in P is tested as part of a sequence which involves also making choices over HEADS and/or (P iff HEADS), then the agent's betting behaviour over P will snap into line and co-ordinate with his

---

[27] That the agent's credence in HEADS would (intuitively) be unchanged by seeing the coin land P side up holds only because the agent's credence in P is the range (0,1). In scenarios where the agent's credence in P at t_0 is not symmetrical, intuitively the agent's credence in HEADS should change (Joyce 2010; Sturgeon 2010; Hart and Titelbaum 2015). See also footnote 29.

[28] Similarly, this holds only because the agent is certain that the coin is fair.

or her choices over the other propositions: in these scenarios the agent will bet as though his or her credence in P was 0.5. To do otherwise would be to act irrationally and expose him- or herself to the risk of being Dutch-booked. Thus the toy model described in the last paragraph represents the way that an agent might reasonably behave in White's situation—and on the dispositional account, such behaviour meets the rationality requirement.[29]

What would the imprecise probabilist say about an agent who behaved as described in the toy example? The imprecise probabilist models an agent with a set of credence functions—so which credence functions belong in this set? It is not usually part of the imprecise probabilist's account to explain exactly how an agent's actual and dispositional choice behaviour determine which credence functions belong in the set that models the agent's credal state. But if we assume that the imprecise probabilist adopts the decision theory 'caprice' (see Sect. 7), then we can draw some conclusions.[30] If the agent is rational, then the agent's choice behaviour in any scenario must maximise expected utility relative to at least one credence function from the set that represents his or her credal state. So given that across the scenarios in which the agent is subjected to De Finetti's test just over P the agent produces the full range of numbers between 0 and 1, it follows, if the agent is rational, that in the set that represents his or her credal state at t_1 there must be credence functions that between them assign all the numbers between 0 and 1 to P. Presumably every credence function in the set assigns 1 to (P iff HEADS), because (as the situation is set up) the agent becomes certain of this at t_1. Therefore if the agent is rational—and so there are only probability functions in the set—each probability function will assign the same number to HEADS as it assigns to P. Thus there must be credence functions in the set that between them assign all the numbers between 0 and 1 to HEADS. This agent then—if he or she is rational and so can be modelled at all—will be modelled as though his or her credence in HEADS has dilated to the range (0,1). There is no way of reading off from this model the fact that the agent will consistently choose as though his or her credence in HEADS is 0.5. The model will then be—at worst—misleading (if we infer from the claim that

---

[29] What about 'good dilation'? Suppose for example that we adjust White's scenario so that at t_0 you have a credence of 0.5 in HEADS and a credence that ranges strictly between 0.6 and 0.8 in P. Then, once the coin is tossed and you see at t_1 that it has landed P side up, intuitively your credence in HEADS and/or P should change: you cannot rationally retain both a credence of 0.5 in HEADS and a credence of (0.6, 0.8) in P. The dispositional account gets the right result here. To see this, consider a counterfactual scenario at which you are tested on all three claims: HEADS, (P iff HEADS), and P. You cannot produce the number 0.5 for HEADS, 1 for (P iff HEADS), and some number between 0.6 and 0.8 for P—at least not if your betting behaviour in this scenario fits with some probability function, as it must if you are rational. It turns out that on the dispositional account whatever the agent's credence in HEADS and P at t_0, if at t_1 the agent becomes certain of (P iff HEADS), then there must be an overlap between the agent's credence at t_1 in HEADS and in P. Many thanks to an audience member at the Bristol Formal Epistemology Gathering for this challenge.

[30] We could draw other different conclusions for versions of imprecise probabilism that include other decision rules. For example—though I don't attempt to prove this here—it follows from the version of imprecise probabilism that includes the decision rule 'maximin', Gärdenfors and Sahlin (1982) and Gilboa and Schmeidler (1989) that an agent who behaved as described in the toy model above would be irrational. The same follows for many of the decision rules that imprecise probabilists have proposed: whether it follows for Robert Williams' rule 'Randomize' (Williams 2014) is an interesting question.

an agent's credence in HEADS is the range (0,1) that his or her actual and dispositional choice behaviour will reflect this full range), or—at best—less informative than it might be. In particular, the imprecise probabilist's model will be less informative than the model that we get on the dispositional account, from which we can see that the agent's betting behaviour over HEADS will be in line with a credence of 0.5 in this proposition.[31]

8 Conclusion

The dispositional account is an alternative to imprecise probabilism. Though on the surface the dispositional account looks less elegant than imprecise probabilism, it is actually simpler because the dispositional betting behaviour that *defines* an agent's credal state also figures in the model that *represents* the agent's credal state. The dispositional account can also produce models that are more informative than those we get with Imprecise Probabilism, as illustrated in the case of dilation.

References

Augustin, T., Coolen, F., de Cooman, G., & Troffaes, M. (2014). *Introduction to imprecise probabilities*. New York: Wiley.
Bradley, R. (2009). Revising incomplete attitudes. *Synthese, 171*(2), 235–256.
Bradley, R. (2017). *Decision theory with a human face*. Cambridge: CUP.
De Finetti, B. (1964). Foresight: its logical laws, its subjective sources. In H. E. Kyburg & H. E. Smokler (Eds.), *Studies in subjective probability* (pp. 93–158). New York: Wiley.
Elga, A. (2010). Subjective probabilities should be sharp. *Philosophers' Imprint, 10*(5), 1–11.
Eriksson, Lina, & Hajek, Alan. (2007). What are degrees of belief? *Studia Logica: An International Journal for Symbolic Logic, 86*(2), 183–213.
Gärdenfors, P., & Sahlin, N.-E. (1982). Unreliable probabilities, risk taking and decision making. *Synthese, 53,* 361–386.

---

[31] This information would also be lost on the model that the dispositional account would produce if we restricted the credence functions that could feature in the model to probability functions. To see this, take as an example the scenario in which the agent is tested just over P and produces the number 0.3. There is no probability function that endorses this choice behaviour and also assigns the number 1 to (P iff HEADS) and 0.5 to HEADS. On this model then, the probability functions that correspond to this scenario must all assign a number other than 0.5 to HEADS or a number other than 1 to (P iff HEADS). Thus we cannot read off from this model that the agent will always choose as though his or her credence in HEADS is 0.5 and his or her credence in (P iff HEADS) is 1.

Gilboa, I., & Schmeidler, D. (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics, 18,* 141–153.

Gillies, D. (2000). *Philosophical theories of probability*. Oxford: Routledge.

Hammond, P. (1988). Orderly decision theory. *Economics and Philosophy, 4,* 292–297.

Hart, C., & Titelbaum, M. G. (2015). Intuitive dilation? *Thought: A Journal of Philosophy, 4*(4), 252–262.

Jeffrey, R. (1983). Bayesianism with a human face. In J. Earman (Ed.), *Testing scientific theories*. (pp.133–156). University of Minnesota Press.

Joyce, J. (2005). How probabilities reflect evidence. *Philosophical Perspectives, 19,* 153–178.

Joyce, J. (2010). A defense of imprecise credences in inference and decision making. *Philosophical Perspectives, 24*(1), 281–323.

Levi, I. (1974). On indeterminate probabilities. *Journal of Philosophy, 71,* 391–418.

Lewis, D. (1973). *Counterfactuals*. London: Blackwell.

Mahtani, A. (2018). Imprecise probabilities and unstable betting behaviour. *Noûs, 52*(1), 69–87.

Pederson, A., & Wheeler, G. (2014). Demystifying dilation. *Erkenntnis, 79*(6), 1305–1342.

Rabinowicz, W. (2014). Safeguards of a disunified mind. *Inquiry, 57*(3), 356–383.

Ramsey, F. P. (1931). Truth and probability. In F. P. Ramsey (Ed.), *The foundations of mathematics and other logical essays* (pp. 156–198). Oxford: Routledge.

Seidenfeld, T., & Wasserman, L. (1993). Dilation for sets of probabilities. *Annals of Statistics, 21,* 1139–1154.

Skyrms, B. (2000). *Choice and chance* (4th ed.). Ontario: Wadsworth.

Sturgeon, S. (2010). Confidence and coarse-grained attitudes. In T. S. Gendler & J. Hawthorne (Eds.), *Oxford studies in epistemology*. Oxford: OUP.

Walley, P. (1991). *Statistical reasoning with imprecise probabilities*. London: Chapman & Hall.

Weatherson, B. (2008). Decision making with imprecise probabilities. *Unpublished Manuscript*. https://philpapers.org/rec/WEADMW.

White, R. (2010). Evidential symmetry and mushy credence. In T. Szabo Gendler & J. Hawthorne (Eds.), *Oxford studies in epistemology*. (pp. 161–186). Oxford University Press.

Williams, R. (2014). Decision-making under indeterminacy. *Philosophers' Imprint, 14*(4), 1–34.