Philosophical Quarterly, forthcoming

Interpreting Imprecise Probabilities

Nicholas J.J. Smith

Abstract

In formal modelling, it is essential that models be supplied with an interpretative story: there must be a clear and coherent account of how the formal model relates to the phenomena it is supposed to model. The traditional representation of degrees of belief as mathematical probabilities comes with a clear and simple interpretative story. This paper argues that the model of degrees of belief as imprecise probabilities (sets of probabilities) lacks a workable interpretation. The standard interpretative story given in the literature is shown to lead to unacceptable results—and, it is argued, there is no way for imprecise probabilists to restrict, replace or finesse this story so as either to avoid or to be able to live with the consequences. Thus the imprecise probabilist lacks a viable account of how a set of probability functions models a belief state: a story about which properties of the model correspond to facts about the thing modelled and how exactly we are to extract information about an agent's belief state from a set of probabilities.

keywords: degree of belief, credence, probability, imprecise probability, formal model

word count: 9,529

1 Introduction

The received formal model of degrees of belief represents them as probabilities. A rational agent's belief state is represented by a probability function over a suitable space or algebra of worlds, events or propositions, with the value assigned to a proposition representing the agent's degree of belief in that proposition. In the philosophical literature, the most prominent alternative to this received view represents degrees of belief as imprecise probabilities:[1] a rational agent's belief state is represented not by a single probability function but by a set of such functions.[2]

[1]Joyce (1998, 600) makes a stronger claim about the status of the imprecise probability approach: that it is not merely the most prominent alternative to the received view but is in fact the standard view: 'Most probabilists recognize that opinions are often too vague to be pinned down in numerical terms, and it has therefore become standard to represent a person's partial beliefs not by some single credence function but by the class of all credence functions consistent with her opinions.' Rinard (2015) also refers to this view as the 'standard fix' for problems facing the traditional view.

[2]In the philosophical literature, the term 'imprecise probability' is generally (although not invariably—cf. n.22) used in connection with a particular kind of model of belief states: one which models them as sets of probability functions. (See e.g. the opening sentence of Bradley and Steele (2014): 'Imprecise probabilism—the view that your belief or credal state is best represented by a *set of probability functions*—has received a lot of attention recently.') In the literature in computer science, artificial intelligence, expert systems, engineering, statistics and elsewhere, the term 'imprecise probability' is used as an umbrella term for non-probabilistic models of uncertainty, in which uncertainty is modelled by

One of the key lines of argument in favour of imprecise probabilities over the received view is that imprecise probabilities are better able to model the belief states of agents who lack detailed evidence bearing on the propositions over which their degrees of belief are defined.[3] One of the key lines of argument in favour of the received view over imprecise probabilities is that while the former comes as part of a package involving a mature and successful decision theory—expected utility theory—there is no plausible decision theory that meshes with the latter.[4] In the ledger of pros and cons, this paper offers a new mark in favour of precise probabilities and against imprecise probabilities.

A feature essential in any formal model is that there be a clear and coherent story about how the models relate to the phenomena they are supposed to represent. To illustrate, consider three different formal representations of colours: as points in RGB space, as points in HSB space, and as strings of hexadecimal digits. When asked what colour a certain formal representation in each system represents and why, we get different answers. In the first case the representation—say 16%, 27%, 7%—comes in three parts, representing the amounts of red, green and blue light (respectively) needed to produce the colour. In the second case the representation—say 92°, 65%, 33%—comes in three parts, but this time they represent the hue, saturation and brightness components (respectively) of the colour. In the third case the representation—say `38551E`—comprises six digits. The first/second/third pair represents the same thing as the red/green/blue component (respectively) in the first representation (this time as a number from 0 to 255, in hexadecimal). The necessity to tell a story about how exactly the model relates to the phenomena is reflected in the individuation of models. Even though one can translate between RGB and HSB, they are regarded as distinct colour models—because there are quite distinct stories about how the representations relate to colours. In the RGB model we can change from one colour to another by increasing the amount of red light or reducing the amount of green, and so on. In the HSB model the story is quite different (even if ultimately intertranslatable): we can change the hue (move across the rainbow) or increase the saturation (make the colour richer), and so on. And for the third representation there is no independent story about how the strings of digits represent colours: the story goes via RGB. For precisely this reason, hex values are not regarded as a distinct colour model, but rather as a particular presentation or encoding of the RGB model (one that is useful for example in computer coding contexts such as website creation in HTML and CSS).

A formal model of degrees of belief likewise requires a story about how we go from a representation within that modelling system, to information about the thing modelled: an agent's degrees of belief. There needs to be a clear and coherent account of the relationship between the models and the phenomena. In the traditional probabilistic model of degrees of belief, the story is simple and direct. A probability function assigns

something other than a single probability function. (See e.g. Walley (2006, 3353): '*Imprecise probability* is used as a generic term to cover all mathematical models which measure chance or uncertainty without sharp numerical probabilities.') There are many such models, including (to name just some, and to cite only a few key works) upper and lower probabilities (Smith 1961), Dempster-Shafer belief functions (Dempster 1967, 1968; Shafer 1976), sets of probability functions (i.e. imprecise probabilities in the first sense noted above) (Levi 1974, 1980; Kaplan 1983, 1996; van Fraassen 1990; Joyce 2010), upper and lower previsions (Walley 1991; Troffaes and de Cooman 2014), and probability intervals (Kyburg and Teng 2001). In this paper I use 'imprecise probability' in the first, philosophical sense.

[3] See e.g. Levi (1985, 396) (cf. also Levi (1974, 394–5)), Kaplan (1996, 27–8 (cf. also 24, 29)), Joyce (2010, 283, 285), Walley (1991, 34, 246 (cf. also 7)) and Sturgeon (2008, 159).

[4] See e.g. Elga (2010), White (2010) and Bradley (2019).

numbers in the interval $[0, 1]$ to propositions (or subsets of the sample space)—and those numbers directly represent strength of belief. 0 means the weakest possible strength of belief. 1 means maximum strength belief. In between, bigger numbers mean stronger degrees of belief and smaller numbers mean weaker degrees of belief.[5]

What about the imprecise probability model? When we take a *set* of assignments of numbers to propositions (or subsets of the sample space) how exactly is that supposed to represent degrees of belief? That is, if you give me a set of probability functions, what information can I draw from it about the belief state of the agent that it represents—and how do I extract this information?[6] The standard story is that descriptions satisfied by all the probability functions in the set—or properties that they all have in common—correspond to facts about the belief state being modelled. As Joyce (2010, 287) puts it:

> A believer's overall credal state can be represented by a *family* C of credence functions defined on some Boolean algebra $\mathbf{\Omega}$. Facts about the person's opinions correspond to properties common to all the credence functions in her credal state.[7]

So the information that a *set* of probabilities conveys about a belief state is the collection of all those pieces of information that *every* probability in the set conveys individually—or would convey, if it alone were taken as a representation of degrees of belief, as in the traditional precise probability model. That is, think of each probability function in the set as a representation of degrees of belief (in the traditional way indicated in the previous paragraph) and then think of the set as a whole as conveying those pieces of information that can be extracted from *every* individual representation. Here the metaphor of a credal committee is sometimes used. As Joyce (2010, 288) puts it:

> Think of C as a huge committee in which each member (= credence function) has a definite degree of confidence for each proposition in $\mathbf{\Omega}$. The person's credal state is a kind of amalgam of the opinions of her committee members, where the amalgamation process is constrained by Pareto considerations: if all members agree about some matter this reflects a determinate fact about what the person believes.[8]

So we think of each probability function in the set as a traditional precise probability model of a (fictional) agent (a committee member). Then we think of the information conveyed by this set about the (real) agent's belief state as the descriptions that are true (according to that traditional model) of the belief states of *all* those fictional agents.

[5]Note that (as in the colour case) we are concerned here with a general story about how the models relate to the phenomena, not with a method of measuring the degrees of belief of a given agent (or the colour of a given sample). Cf. Cohen's (1977, 6) distinction between a criterion of gradation and a method of assessment.

[6]As mentioned in n.5, we are primarily concerned here with the issue of interpreting formal models (in this case sets of probability functions)—with the passage from a formal representation within the modelling system to conclusions about the thing modelled (in this case an agent's belief state)—rather than with the problem of measurement—of how to come up with one particular formal representation within the modelling system, given certain (observable) facts about the thing modelled (although the latter issue will be touched on at certain points, e.g. in the discussion of view (1) in §4.2).

[7]See also Joyce (2005, 156): 'Determinate facts about the person's beliefs correspond to properties that are invariant across all elements of c_t.' Cf. Kaplan (1989, 57).

[8]Cf. Moss (2018, 163).

Here's a concrete way of thinking about the matter. Take possible descriptions of or propositions about the belief state of an agent and write each of them on one row in a ledger that has a column for each probability function in the set. Then for each cell in the ledger—determined by a row with a description or proposition at the left, and a column with a probability function at the top—place a tick in that cell just in case, in the traditional probabilistic model of degrees of belief, that description or proposition would be true of an agent modelled by that (single) probability function. Now the descriptions or propositions that are true of the agent modelled by the whole set of probability functions are exactly those that have a tick in every column.[9]

For ease of reference let's call this account of how claims about—or descriptions of—an agent's belief state are to be extracted from a set of probabilities the 'Pareto' account.

2 Problem

I'll now argue that the Pareto account generates unacceptable consequences. As will be explained further below, these consequences are of two kinds: those that render imprecise probabilism *self*-undermining, and those that render it unacceptable from an *external* common-sense standpoint.

First consider a common kind of evidential situation: one involving nonspecific evidence. One often has evidence that the truth fits some general description without having any evidence that it fits, or does not fit, any more specific description. For example, a detective might discover evidence that lends weight to the view that the perpetrator was female (suppose that a certain parrot was observed to squawk and that it has been noted in the past that the parrot only squawks in the presence of females) but lends no weight to the view that the perpetrator was Apate in particular (likewise Brimo, Circe and all the other possible female perpetrators). For another example, if you get food poisoning, the timing of the onset of symptoms may provide evidence that it was caused by something you ate at a certain restaurant—but it provides no evidence for the view that it was the fish in particular (likewise the chicken, the prawns, the salmon mousse and all the other dishes you sampled). The imprecise probabilist models such situations as follows. Suppose we want to say that the evidence lends probability 0.4 to the claim that a woman did it while lending no particular probability to the claim that Apate (Brimo

[9] Note that there are two different kinds of propositions in play. First, each probability function is defined over a space or algebra of events or propositions—so (depending on the setup) we have propositions directly in the algebra, or propositions that directly describe events in the sample space. Think here of propositions such as 'The die will land on 6', 'The die will not land on an even number', and so on. Second, we have propositions about the belief states of agents. The first kind of proposition—propositions in the algebra—may in certain special cases include propositions about the belief states of agents, but in general will not: they might just be propositions about how a die will land, or about the weather tomorrow, and so on. In the project of formal modelling of the belief states of agents using (sets of) probability functions, we want the model to yield propositions of the second kind—propositions about the belief states of agents (which is, after all, what we are trying to model). For example, when we model Xen's belief state about the outcome of a sporting competition by a probability function (or set of probability functions) over an algebra of propositions such as 'Team A will win', 'Team B or Team C will win', etc., we don't want our model to output one or more of these propositions (the ones in the algebra): we want it to give us propositions of the form 'Xen assigns probability greater than 0.5 to Team A winning', or 'Xen is confident that neither Team B nor Team C will win' and so on. To avoid any possible confusion, let me note explicitly that in what follows, we never suppose that probability functions assign values to propositions of the second kind—propositions about the belief states of agents. Probability functions assign values to propositions of the first kind—propositions in the algebra—and on the basis of modelling an agent by one or more such probability functions, the theorist or modeller makes claims of the second kind—claims about the belief state of the agent modelled.

etc.) did it. The model involves a set of probability functions, all assigning 0.4 to the set of all women (in the sample space of possible perpetrators) and between them assigning all possible values (compatible with their being probability functions) to the singletons {Apate}, {Brimo} etc. But there is a problem here, given the Pareto account. In this case it is correct to describe each committee member as proceeding as if (A) is true:

(A) The evidence lends a precise weight to the claim that Apate is guilty

(they just disagree over what weight it is) and incorrect to describe any committee member as proceeding as if (B) is true:

(B) The evidence lends no particular weight to the claim that Apate is guilty.

According to the Pareto account, claims that are true of all committee members represent facts about the agent's belief state. Yet the imprecise probabilist's aim here is to model an agent for whom (B) is true and (A) is false.

The problem is a more general one. The imprecise probability view is meant to be able to model an agent for whom there are propositions P such that there is no *particular* or *unique* real number x such that the agent believes P to degree x. But each probability function in the set representing the agent assigns a particular number to each proposition.[10] Hence considering each probability function individually, we may conclude the following about the belief state of the committee member that this function models:

(C) There is a unique real number x such that P is believed to degree x.

But different functions in the set assign different values to P. So there is no real number k such that considering every probability function individually, we may conclude the following about the belief state of the committee member that this function models: 'P is believed to degree k'. So given the Pareto account, (C) represents a fact about the agent's belief state, but there is no real number k such that 'P is believed to degree k' represents a fact about the agent's belief state. So far from being a common-sense model of an everyday kind of agent—in contrast to the artificial, unworldly precision of traditional probabilism—this is unacceptable from a common-sense viewpoint. How can it be true that *there is* a unique real number that represents the agent's degree of belief in P *without* there being any real number at all such that *it* represents the agent's degree of belief in P? This is like claiming that there is one green apple in the bowl while at the same time denying of each apple individually that it is green.

Here's another manifestation of the problem. Consider an agent who does not have enough evidence to say, of two particular events H and T, that they are equally likely, or that H is more likely, or that T is. For example, suppose that just prior to tossing an ordinary coin the agent drops it and it gets run over by a tram. It now looks somewhat worse for wear but still fairly flattish and symmetrical. H (T) is the event of its coming up heads (tails) on the next toss. The imprecise probabilist will model the agent's degrees of belief by a set that includes probability functions that assign H and T the same value, functions that assign H a higher value, and functions that assign T a higher value. But every probability function assigns a particular real number in the interval $[0, 1]$ to each

[10] It could be the case that certain (non-measurable) subsets of the sample space are assigned no probability. However this point is not germane here, for imprecise probabilists try to model agents who have no particular real numbered degree of belief in P not by saying that P is non-measurable (a route that would also be open to precise probabilists) but by saying that P is assigned different values by different probability functions in the set representing the agent.

event (for which it is defined—each measurable subset of the sample space). So for any probability function p and any such events X and Y, either $p(X) = p(Y)$ or $p(X) > p(Y)$ or $p(X) < p(Y)$. So given the Pareto account, 1 below represents a fact about the agent's belief state—but none of 2, 3 or 4 does:

1. For any X and Y, X and Y are believed to exactly the same degree, or X is believed more strongly than Y, or Y is believed more strongly than X
2. H and T are believed to exactly the same degree
3. H is believed more strongly than T
4. T is believed more strongly than H.

This is unacceptable as a model of a rational agent. Think of the model as an oracle to which we put a claim about the agent's belief state. Internally, the oracle considers the claim in relation to each committee member individually—in each case employing the traditional precise probability model—and if it notes that the claim is true of each and every committee member, then the oracle returns True (and otherwise Not True). First I put claim 1 and the oracle returns True. 'Great,' I say: 'Well, I'm particularly interested in H and T—so which of 2, 3 and 4 is true?' 'None of them,' says the oracle. 'Wait, what?!' I object: 'But you just said...!' The combination of the responses Not True to claims 2, 3 and 4 makes it the case that claim 1 isn't true. Yet the oracle responds True to claim 1. This is unacceptable as a model of a rational agent's belief state.

The problem is not restricted to properties that hold of all probability functions. For example, it is not the case that for every probability function, there is a proposition to which it assigns the value 0.4. But it could certainly turn out that the set of probabilities that the imprecise probabilist takes to model some agent contains probability functions each of which, as it happens, makes some assignment of 0.4. It could also be the case that there is no one proposition that is assigned 0.4 by all of these probability functions. In this case 'There is some proposition that is believed to degree 0.4' will state a fact about the agent's belief state—but there will be no proposition P such that 'The agent believes P to degree 0.4' states a fact about the agent's belief state. Again this is unacceptable as a model of the belief state of a rational agent.

There are two kinds of problem brought out by the foregoing examples. One problem is that the imprecise probabilist's original motivation is to be able to say things about agents that cannot be said on the traditional precise probabilist view—for example:[11] there are propositions to which the agent assigns no unique, precise probability; there are propositions such that the agent neither assigns them exactly the same probability, nor assigns one a strictly greater probability; there are propositions P such that the agent thinks P's probability is at least 0.5, or thinks P is more likely than $\neg P$, but does not assign P (or $\neg P$) any particular probability; there are propositions P such that the agent thinks P's probability is around 0.5 but not exactly 0.5 or 0.49 or 0.51 or any other precise number; and so on. Yet on the Pareto view, the *negations* of these claims are truths about all agents—because they correspond to facts about probability functions, and hence are true of all committee members. In other words, the above claims are all false descriptions of all committee members—because every committee member

[11] In addition to quotes and references elsewhere in this paper, see e.g. Levi (1974, 395; 1980, 185), Jeffrey (1983, 144; 1987, 586, 587), van Fraassen (1990, 347), Kaplan (1996, 24), Joyce (2005, 156), Rinard (2013, 158; 2017, 259) and Mahtani (2018, 69).

is a precise probabilistic agent. Their negations are true descriptions of all committee members—and so true descriptions of all agents, given the Pareto view. So imprecise probabilism is self-undermining.

The second problem is that the view yields descriptions of agents that sound confused or unacceptable from an everyday, common-sense point of view. The core of the issue is as follows. An individual probability function is an internally coherent representation of a belief state in the sense that if the traditional probabilistic model represents the agent as believing some proposition or other to degree 0.4, it will also represent the agent as believing some particular proposition to degree 0.4; if it represents the agent as either believing P more than Q or Q more than P, it will also represent the agent as believing P more than Q (or Q more than P); and so on. This coherence is lost when we move to a set of probability functions and apply the Pareto account. Now the agent can be represented as believing some proposition or other to degree 0.4, without believing any particular proposition to degree 0.4—and so on across the other manifestations of the problem. The problem lies in a fundamental disconnect between the general or high level representations of the agent's belief state (as believing something of such and such sort, as believing this *or* that, etc.) and the particular or low level representations (as believing this proposition to this degree, that one to that degree, etc.). The former can be true (according to the model) without the latter being such as would be required to make them true. In committee terms, when a high level claim is true of different committee members for *different* low level reasons (but still true of all of them) it will come out as being true of the agent for *no* reason.

I'll now consider responses to the foregoing argument against the Pareto account. The responses fall into four categories (and will be discussed in the sections indicated):

(i) I have misunderstood imprecise probabilism or the Pareto account. (§3)

(ii) We should modify the Pareto account: in particular we should restrict the descriptions that can go into the ledger in the first place in such a way that the problematic results no longer follow. (§4)

(iii) We should replace the Pareto account with a different story about how to interpret sets of probabilities, from which the problematic results do not follow. (§5)

(iv) The results are not actually problematic. (§6)

3 Misunderstanding

The reader may think that I have misunderstood the imprecise probability view, and point to the following passage from Joyce (2010, 288):

> It is sometimes suggested that this model is even more psychologically implausible than its precise cousin since believers must keep track of a (typically infinite) family of credence functions, rather than just one. This is the wrong way to think. Rather than being a model of a believer's psychology, the credal state is a highly formalized representation of her doxastic situation. Though the person's opinions are modeled by the shared properties of her "committee members", she herself will not think in these terms. Instead, she will make qualitative or comparative assessments of probability and utility—that X is more likely than Y, that X and Y are independent, that X is the evidence for Y, that A is a better act than A^*, *et cetera*—and these concrete judgments are

> modeled abstractly by requiring that all $c \in \mathrm{C}$ satisfy certain conditions, e.g., $c(X) > c(Y)$, $c(X\&Y) = c(X) \cdot c(Y)$, $c(Y|X) > c(Y)$, $Exp_c(A) > Exp_c(A^*)$, *et cetera*. The believer only keeps track of her explicitly held qualitative and comparative beliefs: the formal representation takes care of itself.[12]

But Joyce is here responding to a different objection—and his response does not connect with my argument.[13] To illustrate with reference to just one of the examples from §2, my point is that, according to the Pareto account, (C) is true of the agent, because it is true of every committee member. It is irrelevant here whether the agent explicitly keeps track of anything like (C). I am not saying that the Pareto account implies any explicitly held belief on the part of the agent: I am simply saying that it implies that (C) represents a fact about the agent's belief state—and this in itself is problematic, for two reasons. First, it undermines the motivation for imprecise probabilism, which was to be able to say (among other things) that there can be agents of whom (C) is *false* (for some proposition P). Second, it is unacceptable as a model of a rational agent (when combined with the fact that there is no real number k such that 'P is believed to degree k' represents a fact about the agent's belief state). In the terms of the quotation from Joyce, the problem arises because (C) is part of the representation of the agent's doxastic situation—it is not the case that we get a problem only if the agent herself thinks in these terms. Therefore pointing out that the latter is not part of the imprecise probability picture does not help with my objection.

Second, I have encountered the following response to my argument. As Joyce (2010, 293–4) notes: 'with any attempt to represent some phenomenon mathematically, it is critical to figure out which aspects of a representation reflect the reality being modeled and which are artifacts of the formalism.' So—the response goes—one can simply deny that (C) (for example) represents a fact about the agent's belief state and claim instead that it is representationally insignificant: an artefact of the formalism. However, this response is inadequate. The general distinction between artefacts and representational features of a formal model is certainly a legitimate and important one—but noting that this distinction exists is only the beginning. One must then go on to say exactly which features of the model are artefacts and which represent facts about the phenomena being modelled. Joyce makes this point himself[14]—and it had earlier been made very clearly in other contexts, for example regarding formal models of vagueness:

> Of course, saying that the account is meant to be a model, and thus that certain unattractive parts of the semantics are artifactual, is not enough. We have yet to determine in general which aspects of the model are artifacts and which are representors . . . Without knowing in more detail what is representor and what is artifact we cannot draw any useful insights from the model, since we do not know what parts of it are intended to provide such information. (Cook 2002, 240–1)

> it must be determined which features of formal languages correspond to relevant features of correct reasoning in natural language, and which features do

[12]I omitted a comma from the quotation (after $c(X\&Y)$) as it is clearly a typo.

[13]This is an objection not to Joyce but to a (possible) reader who thinks that Joyce has responded to my argument.

[14]Again, I am not objecting to Joyce, but to someone who mistakenly thinks that Joyce has given a response to my argument.

> not. Otherwise, there is a danger of inferring something about the target ... on the basis of an artifact of an otherwise good model. (Shapiro 2006, 50)

The Pareto account is, precisely, the proposed way of drawing the distinction between representors and artefacts:[15] properties common to all probabilities in the set correspond to facts about the agent's belief state. So one cannot simply claim, in response to my objection, that (C) is an artefact: given the Pareto account, (C) *does* represent a fact about the agent's belief state.

4 Restriction

What if we make further restrictions on which aspects of the set of probabilities represent facts about the agent? More specifically, we retain the idea that descriptions that get a tick in every column are true of the agent—but restrict the descriptions that can go in the ledger in the first place. In this section I'll criticise two specific proposals of this sort before giving a general reason why no restriction strategy can avoid all the problems set out above.

4.1 Elementary

One option would be for the imprecise probabilist to say that only elementary facts about each probability function—that it assigns 0.2 to this proposition and 0.5 to that one etc.—correspond to facts about the agent's belief state, as opposed to facts such as 'There is a number x such that P is assigned x'. In other words, the disconnect between high level and low level representations (discussed in §2) is resolved by jettisoning the high level representations altogether. Recall (from n.9) the distinction between propositions in the algebra and propositions about the belief states of agents. The proposal now on the table is that only the most elementary or low level propositions of the latter kind be allowed: those of the form 'The value assigned to P is n'—or 'P is believed to degree n'—where P is in the algebra.[16]

Earlier I noted two kinds of problem with the Pareto account: it renders imprecise probabilism self-undermining, and it yields descriptions that are, from an external common-sense standpoint, unacceptable as descriptions of rational agents. The restriction just proposed is problematic from both these points of view. First, it renders the imprecise probability model unable to model the kinds of things that are often taken to motivate the model in the first place: for example agents who are more confident of one thing than another without having a unique degree of confidence in either. For descriptions such as 'P is believed more firmly than Q' are now forbidden. Thus it is quite appropriate that imprecise probabilists such as Joyce (2005, 156) explicitly do *not* impose any such restriction:

> Such invariant facts can come in a wide variety of forms. It might be invariant across c_t that a certain quantity has a specific expected value, or that a

[15]I here use 'representor' in the way Cook does in the passage quoted above: as a term for those features of a model that represent facts about the thing modelled (in contrast with artefacts). 'Representor' is used in a different sense in the literature on imprecise probabilities: as a term for the set of probability functions that is taken to represent an agent's belief state. This usage derives from van Fraassen (1990, 347): see the quote in §4.2 below. 'Representor' is used in a closely related but subtly different sense by Maher (1993, 21).

[16]The proposal is not that we restrict ourselves to propositions of the first kind: propositions in the algebra. That proposal would be a complete non-starter. The whole point is to provide a formal model of the belief states of agents: so we definitely need at least some propositions of the second kind—that is, propositions about the belief states of agents.

particular distribution of probabilities has a uniform, binomial, normal, or Poisson form, and so on.

Second, the restriction now prevents us from describing agents in all sorts of ways in which we ordinarily do want to describe them—and in which we are allowed to describe them on the traditional precise probabilistic view. It is completely out of line with the way we think of single probability functions as representing belief states in traditional precise probabilism. On a traditional model, representing my belief state by a certain probability function captures not only that I believe (say) S to degree 0.2 and T to degree 0.5, but also that my degree of belief in T is greater than my degree of belief in S, and so on. A probability function just assigns values to propositions in the algebra (once you specify the value assigned to every proposition, you completely specify the probability function). So to put the point somewhat metaphorically, we might say that were a probability function able to speak, the only things it could say would be of the form 'The probability of P is n' (where P is in the algebra). However it can (as it were) show a lot more than it can say—and traditionally, when we model an agent's belief state with a probability function, we also say these further things about the agent. For example, by saying that S has value 0.2, T has value 0.5, and R has value 1, the probability function shows that T has a greater value than S, shows that R has the maximum possible value, and shows that T has the same value as its negation—and when we use the function to model the belief state of an agent, we say of this agent not only that she believes S to degree 0.2, T to degree 0.5, and R to degree 1, but also that she believes T more firmly than S, that she is maximally certain of R, and that she is maximally uncertain about T. The proposed restriction on the descriptions that go into the ledger in the first place is thus far too strict: it prevents us from saying all kinds of ordinary things about agents that we normally want to be able to say about them—and can say both in traditional precise probabilism and on the (original) Pareto view.

4.2 Comparative

Here's a second proposed restriction of the Pareto account:[17] we only allow descriptions of the form 'The probability of P is greater than or equal to the probability of Q' or 'P is believed at least as firmly as Q' (where P and Q are in the algebra). There are two problems with this response. The first is that it isn't any better than the first restriction when it comes to allowing the imprecise probabilist to say the kinds of things that motivated her view in the first place. For example, on this view we would not be able to describe an agent: as believing P *much* more firmly than Q and a *little* less firmly than R (but without assigning P, Q and R any precise probabilities—e.g. 0.8, 0.3 and 0.9 respectively); as lacking a precise degree of belief in some propositions; as having a credence of around (but not exactly) $\frac{1}{2}$ —or a mid-strength degree of belief—in P; as believing P and Q to roughly the same degree; as believing P twice as firmly as Q (but without assigning P and Q any precise probabilities); and so on. The second problem is that this view—where we represent agents' belief states by sets of probabilities, but the only information that we can draw from the set about the agent is a binary relation of comparative probability—is not an attractive position in the space of possible views: anyone motivated to move from standard imprecise probabilism to this view would be better off moving to a third view instead. To make this clear, I'll first set out these three views:

[17] Thanks to Michael Nielsen for suggesting a restriction along these lines.

(1) is standard imprecise probabilism with the (original unrestricted) Pareto account. The goal is to give a formal model of the belief state of an agent, and certain facts about the agent constitute our starting point: for example, he believes to degree 0.5 that this coin will land heads; he believes very strongly that it will rain tomorrow; he believes more firmly that Labour will win the next UK general election than that the Democrats will win the next US congressional elections; and so on. Now this motley bunch of miscellaneous facts cannot itself be taken as a formal model of the agent—it's not a coherent formal model at all. So what we do is take *all* the probability functions that are consistent with those facts, and represent the agent's belief state by that whole set. As van Fraassen (1990, 347) puts it:[18]

> a finite and even small number of judgements may convey all there is to our opinion. But then there is a large class of probability functions which satisfy just those judgements, hence which are compatible with the person's state of opinion. Call that his or her *representor* (*class*).

An essential question is then how we are to conclude *other* things about the agent, apart from the facts that determined which set of probabilities was the relevant one in the first place—and the answer is the Pareto account.

(2) is like (1) except that the Pareto account is restricted in the way under discussion in this subsection: to comparative statements.

(3) is a fundamentally comparative approach to probability. On this view, there is something of primary significance or importance about comparative probability (e.g. that it has a very clear behavioural meaning: if one thinks that P is more probable than Q, this means that if one has a fixed amount to stake, one would rather stake it on P than on Q). So the primary facts about agents are facts about comparative probability. However it can be shown that if, in a given agent, the comparative probability relation satisfies certain constraints (and possibly if the space of events or propositions over which the relation is defined has a certain richness), then one can represent the relation by an assignment of *numerical* probabilities to events or propositions (i.e. by a probability function—see e.g. de Finetti (1980); Savage (1972)), or (given weaker constraints) by a *set* of such assignments (i.e. by a set of probability functions—see e.g. Nehring (2009); Alon and Lehrer (2014)). These representation results are interesting and mathematically useful—but note that this view is quite different from imprecise probabilism. For the imprecise probabilist, the fundamental formal model of the agent is a set of probability functions. For the comparative probabilist, the fundamental formal model is a binary relation (of comparative probability): the set of probabilities is then a representation of this relation (in the sense of 'represent' involved in the term 'representation theorem')—and one that is available only when certain constraints are satisfied. Modelling a belief state by a binary relation and then, if certain conditions are met, representing that relation by a set of probabilities is not the same as modelling a belief state by a set of probabilities.

We come now to the point that anyone who accepts my argument against (1) and proposes (2) in response should actually move to (3) instead. To see this, we ask the proponents of (2): where does the set of probabilities that, in your view, is the formal model of the agent's belief state come from in the first place? Suppose they say: just as in (1), it is the set of probabilities consistent with all the initial miscellaneous facts about the agent. In this case their view has a severe internal tension, if not outright inconsistency:

[18] For a more detailed discussion of this—and other—ways of motivating imprecise probabilism see Smith (2022, §6).

they take certain claims about agents as initial data that determine a set of probabilities that models the agent, but then their account of how we extract claims about the agent from the set of probabilities does not allow us to conclude many of these very claims. For example, part of what makes this set of probabilities the correct one to model that agent might be that she believes very strongly (but not to some specific numerical degree) that it will rain tomorrow—but this description of the agent is forbidden on view (2). So the problem is that the restricted Pareto account does not allow one to draw out of the formal model (a set of probabilities) facts that went into determining the model (why it was that particular set of probabilities) in the first place. To avoid this problem the proponents of (2) must say that only comparative probability facts (the agent believes P more firmly than Q, etc.) go into determining the set of probabilities. But then they should adopt view (3) instead of (2). On view (3), the fundamental formal model of the agent is a binary relation—which is simpler than a set of probabilities—and on the version of (2) that we have now arrived at, the binary relation captures *all* the facts that determine the set of probabilities (that, on (2), is the fundamental formal model of the agent) and *all* the facts that we can conclude about the agent from the set of probabilities. Furthermore, the set of probabilities is available only when the comparative probability relation that determines it (and possibly the space over which the relation is defined) satisfies certain constraints. The binary relation itself, however, is always available—even (for example) for an agent with a hugely limited psychology who thinks only that some event A is more probable than another event B (and that's the only opinion he has about anything). There is thus every reason to take the binary relation—rather than the set of probabilities—as the fundamental formal model of the agent's belief state. In other words there is every reason to move from (2) to (3): from imprecise probabilism to comparative probabilism.

I have just argued that an imprecise probabilist who wants to move to (2) should move to (3) instead. I have said nothing against (3). So who really wins the argument here—do I, or does the imprecise probabilist? I do. I have said nothing against (3) because it is not my target in this paper. As discussed above, (3) is not imprecise probabilism—and it does not allow one to say many of the things that motivated imprecise probabilists in the first place. It is far from the position against which I am arguing here: the imprecise probabilist view that is widespread—or even standard (see n.1)—in the philosophical literature. Imprecise probabilists such as Joyce (2010, 285ff.) explicitly contrast their view with comparative probabilism—and argue for the superiority of the former. If the upshot is that my opponent moves from the view I am arguing against to a completely different view that is not at all widespread in philosophy—and that leading proponents of imprecise probabilism reject—then that is a victory for my argument here.

4.3 Problem

I've considered two examples of restrictions on the Pareto account. I'll now point out a general problem for any restriction strategy. In §2 I noted that imprecise probabilism is self-undermining: the motivation is to be able to say (for example) that there are some agents and some propositions P such that there is no particular real number x such that the agent believes P to degree x—but one upshot of the view is that for every agent and every proposition P, 'There is a unique real number x such that P is believed to degree x' is a true description of the agent. Now a restriction strategy might stop the latter being an output of the view—but it cannot make the claim that motivated the view in the first place an output of the view. That is, the strategy might stop us saying the negation of what we wanted to say—but it still cannot allow us to say the thing we did want to say. For it is true of every probability function that it assigns a unique real number to every proposition in the algebra. So if the basic idea is that *only some* descriptions true of all committee members (each of whom is modelled in the traditional way by a probability function) are true of the agent (modelled by a set of probability functions) then the view will never allow us to describe an agent using a claim that is *false* of all committee members—and the motivating statement 'There is no particular real number x such that the agent believes P to degree x' is just such a claim.

5 Replacement

If we cannot fix the Pareto account by restricting the descriptions that can go into the ledger in the first place, maybe we can fix imprecise probabilism by replacing the Pareto account with a different story about how to go from a set of probabilities to claims about the belief state of the agent modelled by that set. I can think of three possible sorts of replacement:

5.1 Some

On the Pareto account, we can apply to the agent descriptions that are true of *all* committee members (that get a tick in every column). What if we replace 'all' with 'some'? On this alternative view, we can apply to the agent descriptions that are true of some committee member(s) (that get a tick in at least one column).

Everything that is true of all committee members is true of some committee member, so we can still say everything we could say before—including all the problematic claims such as (C). However perhaps we have made some progress—because now we *can* also say of the agent 'P is believed to degree k' for some particular k. That is, while we have not avoided the self-undermining problem, haven't we at least avoided the problem of saying things that are unacceptable from a common-sense standpoint? No, not at all. On this new view, we must say things that are, if anything, worse. Consider for example the model in §2 that involves a set of probability functions, all assigning 0.4 to the set of all women and between them assigning all possible values to the singletons {Apate}, {Brimo} etc. In this case, on the new view, we must say of the agent that she believes to degree k that Apate (say) did it, for every possible k: she believes this to degree 0.3, and she also believes it to degree 0 and degree 1 and degree 0.5 and so on. Furthermore, as well as saying that the agent believes to degree 0.7 that Apate did it, we must also say that she does not believe to degree 0.7 that Apate did it—because the latter is true of some committee member. In addition, while we must say that she believes to degree 0.7 that Apate did it, and we must say that she believes to degree 0.9 that Apate did it, we cannot say 'the agent believes to degree 0.7 and to degree 0.9 that Apate did it', because this description is true of *no* committee member. So again, this is unacceptable as a model of a rational agent's belief state.

5.2 Super

The second proposal is that we take all the probability functions in the set and form from them a 'super function' f. Like each probability function, f is a function from the algebra (over which all the probability functions are defined) to [0,1]. It is defined by setting $f(P) = x$ iff every probability function assigns x to P. Now we take this single function f (which is not in general a *probability* function) and use it to model the agent in just the way that a single probability function is taken to model an agent on the traditional view.

The problem with this proposal is that we won't be able to say the kinds of things that motivate imprecise probabilism in the first place—for example 'The agent believes P more than Q (but does not believe either of them to a particular precise degree)'. For if the agent does not believe P or Q to a particular degree, that means that f does not assign either of them a value—and that means we cannot say the agent believes P more than Q. Or to put it the other way around, we can only say that the agent believes P more than Q if f assigns P a higher value than it assigns Q—i.e. it assigns m to P and n to Q where $m > n$. But in that case the agent believes P to degree m precisely and Q to degree n precisely (and note that this can only be the case if *every* probability function assigns m to P and n to Q).

5.3 Extrema

The previous two proposals retained recognisable vestiges of the Pareto account; the final proposal discards it altogether. Given a set of probability functions, for any proposition P in the algebra of those functions we can (the proposal goes) conclude only two things about the agent: her upper probability for P (the supremum of the values assigned to P by functions in the set) and her lower probability for P (the infimum of those values).[19]

One initial issue with this proposal is that it raises a further question of interpretation: what are (the meanings of) upper and lower probabilities? However as mentioned in n.2, there is a literature on upper and lower probabilities—and on related ideas such as upper and lower previsions and probability intervals—and responses can be given to this question. For example, one possible approach would be to follow Walley (1991) (given certain assumptions such as convexity of the set of probabilities) and offer an interpretation in terms of maximum buying and minimum selling prices of gambles.

A more pressing issue arises when we ask the proponent of this view: where does the set of probabilities that, in your view, is the formal model of the agent's belief state come from in the first place? Suppose they say: just as in the original imprecise probability view (as spelled out under (1) in §4.2) it is the set of probabilities consistent with all the initial miscellaneous facts about the agent. In this case their view has a severe internal tension, if not outright inconsistency:[20] they take certain claims about agents as initial data that determine a set of probabilities that models the agent, but then their account of how we extract claims about the agent from the set of probabilities does not allow us to conclude many of these very claims. For example, part of what makes this set of probabilities the correct one to model that agent might be that he believes more firmly that Labour will win the next UK general election than that the Democrats will win the next US congressional elections—but this description of the agent is forbidden on the present view (which only allows us to state his lower and upper probabilities for the individual propositions 'Labour will win the next UK general election' and 'the Democrats will win the next US congressional elections').[21] So the problem is that the proposed view does not allow one to draw out of the formal model (a set of probabilities) facts that went into

[19]Cf. the view called **LP** by Joyce (2010, 294).

[20]Cf. the discussion in §4.2.

[21]It is important here not to confuse the present view with a view on which the formal model of the agent is an assignment of a pair of values (a lower and upper probability) to each proposition. (Cf. the distinction drawn between **LP** and Kyburg's view by Joyce (2010, 294).) On *that* sort of view we might well be able to come up with an account of when the agent believes P *more* than Q (for example: when her lower probability for P exceeds her upper probability for Q). But on the present view, the formal model of the agent is a set of probability functions, and the whole point is that the *only* information we may conclude from it about the agent is, for each proposition, a lower and an upper probability.

determining the model (why it was that particular set of probabilities) in the first place. To avoid this problem the proponent of the view must say that only facts about upper and lower probabilities go into determining the set of probability functions—but the problem here is that different sets of probability functions (over a given algebra of propositions) can determine the very same upper and lower probabilities (for all propositions in the algebra): so specifying upper and lower probabilities for every proposition does *not* determine a particular set of probability functions.

I have just argued against a version of imprecise probabilism (i.e. where the formal model of an agent is a set of probabilities) that replaces the Pareto account with the view that we can only extract upper and lower probabilities for propositions from the set of probability functions. I have not argued against views that (do not involve a set of probability functions and instead) model an agent's belief state with a function that assigns pairs—or intervals, as in Kyburg (1961) and Kyburg and Teng (2001)—of values to propositions. Such views are far from my target in this paper, which is the imprecise probabilist view that is widespread—or even standard (see n.1)—in the philosophical literature. See for example Levi (1980, §9.8) and Joyce (2010, 294–6) for detailed discussion of the differences between their own (imprecise probabilist) views and positions of these other sorts—and arguments for the superiority of their own views.[22] Recall the position at the end of §4.2: if at this point my opponent decides to jump ship from the view I am arguing against to a completely different view that is not widespread in philosophy—and that leading proponents of imprecise probabilism explicitly reject—then that is a victory for my argument here.

6 Acceptance

I argued that the Pareto account faces two kinds of problem. On the one hand imprecise probabilists are motivated by wanting to describe agents in ways that traditional precise probabilism cannot, yet the Pareto view does not yield such descriptions—indeed it yields their denials. On the other hand the Pareto account yields claims that are, from a common-sense standpoint, unacceptable as descriptions of rational agents. I have looked at possible modifications of and replacements for the Pareto account that might avoid the problematic outputs. I turn now to the final response which is to say that the outputs are actually OK. Note immediately that this kind of response cannot help with the first sort of problem. It simply is not acceptable to say: 'Here's the sort of thing I want to be able to say about agents—things that cannot be said on the traditional view... Now my view doesn't actually allow me to say these things—and indeed issues in their denials... But that's OK!' But perhaps the strategy can at least help with the second sort of problem—and that might be some progress, rather than none. I'll now consider, and reject, two ways of arguing that the outputs of the Pareto account are actually acceptable from a common-sense standpoint.

6.1 Ignorance

I claimed that it is problematic to say that (C) 'There is a unique real number x such that P is believed to degree x' represents a fact about an agent's belief state while at the same time, for each real number k, denying that 'P is believed to degree k' represents a fact about the agent's belief state; that it is problematic to say that 'There is some

[22]Note also that while Isaacs et al. (2022) use the *term* 'imprecise probability' in association with a view that models credences using single functions whose values are intervals or pairs of real numbers—rather than sets of functions whose values are single real numbers—they explicitly *distinguish* this view from the 'most common' view that 'ascribes sets of probability functions to agents' (p.902).

proposition that is believed to degree 0.4' states a fact about the agent's belief state while denying, for each proposition P, that 'The agent believes P to degree 0.4' states a fact about the agent's belief state; and so on. I drew an analogy with claiming that there is one green apple in the bowl while denying of each apple individually that it is green. Now someone might point out that the latter sort of thing *is* OK—in cases involving ignorance or lack of information. It is fine for a secret agent to be sure that there is a spy in the room while at the same time not thinking of any individual in the room that *she* (in particular) is a spy; it is fine for a gambler to think that exactly one ticket will win the lottery and not think that it is this one that will win, or this one, and so on through all of them (considered individually); and so on. So what's the problem in the imprecise probability case?

Well, in these cases of ignorance, we take an existential claim to be known—and hence true—without any of its instances being *known*. We don't think it can be true without any of its instances being *true*. Yet the latter is what the imprecise probabilist needs. She is not trying to model an agent who in fact harbours a quite definite degree of belief but does not *know* what it is: she is trying to model an agent who *has no* definite degree of belief. If we translate the structure found in ordinary cases of ignorance over to the case of an agent's belief state modelled by a set of probabilities, we do not get the standard imprecise probability view. We get a very different view: the 'black box' interpretation of sets of probabilities.[23] In this approach, the agent's belief state is taken to *be* a single probability function—but we do not know which one it is, and so we work with a set of probabilities that includes all those that might for all we know be the one 'in the black box'. The set of probabilities models ignorance: it represents a set of epistemic possibilities regarding the agent's belief state. The model of the belief state itself is a single (albeit unknown) probability function. This approach does not face the problems that I have raised: in *this* context, the Pareto account makes perfect sense. That is, if we take what is *true* of the agent to be given by a single *unknown* probability function (hidden in a black box), and we consider the set of all probability functions that might for all we know be the one in the box, and then we take what we *know* about the agent to be given by the Pareto account—i.e. we know a claim to be true of the agent if it is true of all committee members, each of whom is modelled (in the classical way) by a single probability function in the set—then there is no problem. Indeed, this just follows the pattern of the standard treatment of knowledge in epistemic logic (i.e. known = true in all accessible epistemic possibilities; see e.g. Fagin et al. (1995, 15)). This view may have serious problems—e.g. what on earth could *determine* which is the one true probability function?—but the Pareto account (as an account of what we *know* about the agent) is not one of them. However, this view is a version of *precise* probabilism: it takes the belief state to be a single probability function. Imprecise probabilists such as Levi (1980, §9.2; 1985; 2009, 367), Jeffrey (1983, 155) and Kaplan (1989) explicitly contrast their views with this black box approach.

Summing up: the original imprecise probabilist view—including the Pareto account—involves existentials being true while none of their instances is true.[24] That is unacceptable from a common-sense point of view. Ignorance is of no help in making it seem OK—for ignorance involves existentials being true while none of their instances is known (but one is true). That combination is acceptable—but it takes us to the 'black box' view—and that view is far from my target in this paper, which is the imprecise proba-

[23]The term derives from Good (1962).

[24]Analogous comments apply to disjunctions and their disjuncts.

bilist view that is widespread or standard (see n.1) in the philosophical literature.

6.2 Supervaluation

It might now be thought that another familiar view *does* involve exactly this combination of true existentials with no true instances: supervaluationism about vagueness. If the combination is fine there, then why should it be a problem in the context of imprecise probabilism? Indeed, analogies have been drawn between imprecise probabilism and supervaluationism.[25]

Well, actually the combination is not just 'fine' in the context of supervaluationism. In Smith (2008, 84) I call it the 'problem of missing witnesses' and describe it as 'The major objection unique to supervaluationism'.[26] So the combination is problematic there too—but in the minds of many philosophers (i.e. those who ultimately accept supervaluationism about vagueness) the problem is outweighed by other advantages. The dialectic here is actually somewhat intricate. In the vagueness debate, the sorites paradox occupies a central position. Because it is a *paradox*—intuitively obvious premises lead via intuitively correct reasoning to an intuitively unacceptable conclusion—we know from the beginning that every theory of vagueness must involve some unintuitive element(s). At the same time, classical logic is at stake: some prominent approaches to vagueness involve moving, for example, to non-classical many-valued logics. Now while everyone is going to have to say something unintuitive, many philosophers regard abandoning classical logic as the *most* unintuitive thing one can do—they think it should be a last resort. Among views about vagueness that retain classical logic, the two main contenders are supervaluationism and epistemicism. On a cost/benefit analysis, Williamson (1994) thinks that epistemicism comes out ahead, while Keefe (2000) thinks that supervaluationism does—and arguably Keefe has the numbers among the broader philosophical community.[27] So an accurate perspective is not that the combination of true existentials with no true instances is fine in the context of supervaluationism about vagueness: rather, the combination is problematic—but nevertheless supervaluationism about vagueness is widely accepted because its benefits (are felt by those who adopt it to) outweigh its disadvantages.

Now the literature on formal models of degrees of belief is not the vagueness literature. The playing field is quite different. It is not dominated by a central paradox: not the sorites, nor any other.[28] And classical logic is not at stake. So the combination of true existentials with no true instances is, as it always was, a bitter pill—but this time we have no reason to swallow it. When it comes to vagueness, everyone must pay a price—and this combination is (many think) lower than the price of non-supervaluationist views. But in the literature on formal models of degrees of belief it just becomes a cost, with

[25]See e.g. van Fraassen (1990, 347; 2005, 28; 2006, 483), Hájek (2003, 277–8), Rinard (2013, 158; 2015, §1), Mahtani (2019) and Smith (2022, §6). Note that if imprecise probabilism is analogous to supervaluationism, then the view presented in §5.1 is analogous to subvaluationism (on which see e.g. Hyde (1997)).

[26]Cf. also Keefe (2000, 183).

[27]Of course there is much to be said both about the lessons of such surveys in general, and in particular about the connection between the wording of the vagueness question and the debate between supervaluationism and epistemicism—but for what it's worth, in the 2020 PhilPapers Survey (`survey2020.philpeople.org`) the most popular view on vagueness was that it is 'Semantic' (52%); 'Epistemic' was significantly less popular (24%).

[28]Mahtani (2019) presents a view on which the four-place predicate '_at_has a credence in_of_' is vague and is handled by a supervaluationist treatment of vagueness—but she notes that this predicate is *not* sorites susceptible and hence calls it an 'atypical' vague predicate (p.15).

no countervailing benefit. It may be OK to say this problematic thing if everyone has to say some problematic thing: but it not OK to say it for no good reason. It just becomes a reason—as I have been arguing—not to adopt imprecise probabilism.

I can imagine someone objecting at this point: we have to accept that there can be true existentials with no true instances in order to get an adequate account of vagueness—and once we have done so, we can accept this combination in other contexts too at no extra cost. But this is a mistake. Supervaluationists about vagueness think that claims such as 'there is a man in this sorites series such that he is tall and the very next man is not tall' are true, without there being any man m in the series such that 'm is tall and the man next to m is not tall' is true. Accepting this does *not* make them think that a secret agent's claim that there is a spy in the room can be true without there being any p in the room such that 'p is a spy' is true. We may have to accept the problematic combination of true existentials with no true instances in certain cases involving vague predicates—but this does *not* mean that we can just accept it as fine in other contexts as well.[29] In other contexts, the combination is as problematic as ever.

7 Conclusion

The Pareto account yields consequences that both conflict with the motivations for imprecise probabilism and are unacceptable in themselves. There is no way for imprecise probabilists to restrict or replace the Pareto account so as to avoid the consequences and no way for them to accept the consequences as OK. The imprecise probabilist model of belief states as sets of probability functions therefore faces the severe problem of lacking an adequate interpretative story: an account of how a set of probability functions models a belief state—of which properties of the model correspond to facts about the thing modelled and how exactly we are to extract information about an agent's belief state from a set of probabilities.[30]

References

S. Alon and E. Lehrer. Subjective multi-prior probability: A representation of a partial likelihood relation. *Journal of Economic Theory*, 151:476–92, 2014.

S. Bradley. A counterexample to three imprecise decision theories. *Theoria*, 85:18–30, 2019.

S. Bradley and K. Steele. Uncertainty, learning, and the "problem" of dilation. *Erkenntnis*, 79(6):1287–1303, 2014.

L. J. Cohen. *The Probable and the Provable*. Clarendon Press, Oxford, 1977.

R. T. Cook. Vagueness and mathematical precision. *Mind*, 111:225–47, 2002.

B. de Finetti. Foresight: Its logical laws, Its subjective sources (1937). In H. E. K. Jr. and H. E. Smokler, editors, *Studies in Subjective Probability*, pages 53–118. Robert E. Krieger Publishing Company, Huntington NY, 1980.

A. P. Dempster. Upper and lower probabilities induced by a multivalued mapping. *The Annals of Mathematical Statistics*, 38(2):325–39, 1967.

A. P. Dempster. A generalization of Bayesian inference. *Journal of the Royal Statistical Society, Series B (Methodological)*, 30:205–47, 1968.

A. Elga. Subjective probabilities should be sharp. *Philosophers' Imprint*, 10(5):1–11, 2010.

[29] Cf. Keefe (2000, 182).

[30] Thanks to Michael Nielsen and the anonymous referees for very helpful comments.

R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi. *Reasoning about Knowledge.* MIT Press, Cambridge MA, 1995.

I. Good. Subjective probability as the measure of a non-measurable set. In E. Nagel, P. Suppes, and A. Tarski, editors, *Logic, Methodology and Philosophy of Science: Proceedings of the 1960 International Congress*, pages 319–29. Stanford University Press, Stanford CA, 1962.

A. Hájek. What conditional probability could not be. *Synthese*, 137:273–323, 2003.

D. Hyde. From heaps and gaps to heaps of gluts. *Mind*, 108:641–60, 1997.

Y. Isaacs, A. Hájek, and J. Hawthorne. Non-measurability, imprecise credences, and imprecise chances. *Mind*, 131(523):892–916, 2022.

R. Jeffrey. Bayesianism with a human face. In J. Earman, editor, *Testing Scientific Theories*, volume X of *Minnesota Studies in the Philosophy of Science*, pages 133–56. University of Minnesota Press, Minneapolis, 1983.

R. Jeffrey. Indefinite probability judgment: A reply to Levi. *Philosophy of Science*, 54: 586–91, 1987.

J. M. Joyce. A nonpragmatic vindication of probabilism. *Philosophy of Science*, 65: 575–603, 1998.

J. M. Joyce. How probabilities reflect evidence. *Philosophical Perspectives*, 19:153–78, 2005.

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

M. Kaplan. Decision theory as philosophy. *Philosophy of Science*, 50:549–77, 1983.

M. Kaplan. Bayesianism without the black box. *Philosophy of Science*, 56:48–69, 1989.

M. Kaplan. *Decision Theory as Philosophy.* Cambridge University Press, Cambridge, 1996.

R. Keefe. *Theories of Vagueness.* Cambridge University Press, Cambridge, 2000.

H. E. Kyburg, Jr. *Probability and the Logic of Rational Belief.* Wesleyan University Press, Middletown, 1961.

H. E. Kyburg, Jr. and C. M. Teng. *Uncertain Inference.* Cambridge University Press, Cambridge, 2001.

I. Levi. On indeterminate probabilities. *The Journal of Philosophy*, 71(13):391–418, 1974.

I. Levi. *The Enterprise of Knowledge: An Essay on Knowledge, Credal Probability, and Chance.* MIT Press, Cambridge, MA, 1980.

I. Levi. Imprecision and indeterminacy in probability judgment. *Philosophy of Science*, 52(3):390–409, 1985.

I. Levi. Why indeterminate probability is rational. *Journal of Applied Logic*, 7:364–76, 2009.

P. Maher. *Betting on Theories.* Cambridge University Press, 1993.

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

A. Mahtani. Vagueness and imprecise credence. In R. Dietz, editor, *Vagueness and Rationality in Language Use and Cognition*, Language, Cognition and Mind, volume 5, pages 7–30. Springer, 2019.

S. Moss. *Probabilistic Knowledge.* Oxford University Press, Oxford, 2018.

K. Nehring. Imprecise probabilistic beliefs as a context for decision-making under ambiguity. *Journal of Economic Theory*, 144(3):1054–91, 2009.

S. Rinard. Against radical credal imprecision. *Thought*, 15:157–65, 2013.

S. Rinard. A decision theory for imprecise probabilities. *Philosophers' Imprint*, 15(7):

1–16, 2015.

S. Rinard. Imprecise probability and higher order vagueness. *Res Philosophica*, 94:257–73, 2017.

L. J. Savage. *The Foundations of Statistics*. Dover Publications, New York, second revised edition, 1972.

G. Shafer. *A Mathematical Theory of Evidence*. Princeton University Press, Princeton, 1976.

S. Shapiro. *Vagueness in Context*. Clarendon Press, Oxford, 2006.

C. A. Smith. Consistency in statistical inference and decision. *Journal of the Royal Statistical Society, Series B (Methodological)*, 23(1):1–25, 1961.

N. J. J. Smith. *Vagueness and Degrees of Truth*. Oxford University Press, Oxford, 2008.

N. J. J. Smith. Respecting evidence: belief functions not imprecise probabilities. *Synthese*, 200(475):1–30, 2022.

S. Sturgeon. Reason and the grain of belief. *Noûs*, 42(1):139–165, 2008.

M. C. Troffaes and G. de Cooman. *Lower Previsions*. Wiley, Chichester, 2014.

B. C. van Fraassen. Figures in a probability landscape. In J. M. Dunn and A. Gupta, editors, *Truth or Consequences: Essays in Honor of Nuel Belnap*, pages 345–56. Kluwer Academic Publishers, Dordrecht, 1990.

B. C. van Fraassen. Conditionalizing on violated Bell's inequalities. *Analysis*, 65(1): 27–32, 2005.

B. C. van Fraassen. Vague expectation value loss. *Philosophical Studies*, 127:483–91, 2006.

P. Walley. *Statistical Reasoning with Imprecise Probabilities*. Chapman and Hall, London, 1991.

P. Walley. Imprecise probabilities. In S. Kotz, N. Balakrishnan, C. Read, and B. Vidakovic, editors, *Encyclopedia of Statistical Sciences*, volume 5, pages 3353–6. John Wiley & Sons, Hoboken NJ, second edition, 2006.

R. White. Evidential symmetry and mushy credence. In T. S. Gendler and J. Hawthorne, editors, *Oxford Studies in Epistemology*, volume 3, pages 161–186. Oxford University Press, New York, 2010.

T. Williamson. *Vagueness*. Routledge, London, 1994.

The University of Sydney, Australia