

Vagueness and Imprecise Credence¹

Abstract

In this paper I investigate an alternative to imprecise probabilism. Imprecise probabilism is a popular revision of orthodox Bayesianism: while the orthodox Bayesian claims that a rational agent's belief-state can be represented by a single credence function, the imprecise probabilist claims instead that a rational agent's belief-state can be represented by a set of such functions. The alternative that I put forward in this paper is to claim that the expression 'credence' is vague, and then apply the theory of supervaluationism to sentences containing this expression. This gives us a viable alternative to imprecise probabilism, and I end by comparing the two accounts. I show that supervaluationism has a simpler way of handling sentences relating the belief-states of two different people, or of the same person at two different times; that both accounts may have the resources to develop plausible decision theories; and finally that the supervaluationist can accommodate higher-order vagueness in a way that is not available to the imprecise probabilist.

¹ Many thanks to audiences at the 5th LSE conference on Philosophy of Probability (2016), the Experience and Updating workshop in Bochum (2016), the Visiting Speakers Seminar Series in Stirling (2017), and the Knowledge, Belief, Evidence Conference at Oxford (2018), as well as to all the people who gave me invaluable feedback on this work. I'm also very grateful to the Leverhulme Trust, who supported this research.

Introduction

On the orthodox Bayesian account, every rational agent has a precise credence (or degree of belief) in every proposition that she entertains. Many have objected that this claim is implausible. A rational agent does have a precise credence in *some* propositions: for example, if I am about to toss a coin that you know to be fair, then your credence that (HEADS) it will land heads is presumably exactly 0.5. But now consider the proposition (SARDINES) that my neighbour has at least one tin of sardines in her kitchen cupboard. What is your credence in SARDINES? There are reasons to think that you don't have any precise credence in this proposition.

One reason is that nobody knows what your credence is in SARDINES — not even you. If you are asked what your credence is, then it is likely that no particular number will spring to mind. If pushed, you may be able to produce a number, but the number that you produce will be arbitrary. You might say, for example, '0.352', but you could just as easily have said '0.353'. You do not *know* what your credence is in this claim. And it seems odd that you might have a precise credence — and so be in some particular mental state — without knowing what it is.²

Another reason is this: it is not clear why your credence in SARDINES is some particular value (0.352, say), rather than some other nearby value (such as 0.353). What is it about you that makes it the case that your credence in this proposition is exactly 0.352?³

² At any rate, this idea seems odd at first, though an externalist about knowledge may be easily reconciled to it. Thanks to an anonymous referee for pointing this out.

³ There may also be other reasons to doubt that you do or should have a precise credence in SARDINES. For example, James Joyce would argue that your evidence does not justify any particular credence (Joyce, 2010).

Here we might think that both of these problems can be easily dealt with, for on the orthodox Bayesian view, there is a tight relationship between an agent's credence function and her dispositional betting behaviour. Here is Bruno De Finetti on the subject:

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 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. 101–2)

To illustrate De Finetti's method here, we can apply it to elicit your credence in HEADS. We suppose that you are forced to bet with a bookie over HEADS, and the way the bet works is that you give the bookie some sum pS , and in exchange you will get the sum S if and only if HEADS is true. Before the bet is settled, you get to name the rate p (your 'betting quotient'), and then the bookie gets to fix the sum S . The bookie can fix this sum as either negative or positive, and the idea is that this forces you to produce as the rate p your true credence in HEADS. To see why this is, suppose first that you give some high figure such as 0.8 for rate p . Then the bookie will set the sum S as a positive value — let's say as £10. Then you give the bookie £8 (for this is $0.8(\text{£}10)$), and the bookie will give you £10 back if and only if HEADS is true. Thus you are left committed to a bet which is — by your own lights — a bad deal for you. Now suppose instead that you give some low figure such as 0.2 for rate p . Then the bookie

will set the sum S as a negative value — let's say — £10. Then the bookie will give you £2 (for this is equivalent to your giving the bookie $0.2(-£10)$), and you will have to give the bookie £10 if and only if HEADS is true. Again, you are left committed to a bet which is a bad deal for you. The only way to ensure a neutral deal is to set p equal to your own credence in HEADS — i.e. 0.5.

Can we similarly use this method to elicit your credence in SARDINES? Presumably if we elicited your betting quotient as De Finetti recommends then you would manage to produce some rate p , but the number that you produce would be arbitrary. You would have no good reason to choose the number 0.352, say, over 0.353. Forced to pick a particular number, you might decide on a whim, or choose at random. We can see this clearly by thinking about the betting quotients that you would produce across close possible worlds, where you have the same evidence and rationality as in the actual world. If your betting quotient is elicited in several of these worlds, then the answers you give across these worlds will vary. For you are just deciding randomly or on a whim, and the results of these random or whimsical processes will vary across close possible worlds. We can contrast this with the case where we elicit your betting quotient for HEADS: presumably in each close possible world where you have the same evidence and rationality as you have in the actual world, you will produce the very same number (0.5) when your betting quotient for HEADS is elicited. In this way your betting behaviour across close possible worlds is *stable* where HEADS is concerned, but *unstable* where SARDINES is concerned. I have argued elsewhere (Mahtani, 2016) that this sort of instability in betting behaviour is typical of the sorts of cases that motivate theorists to resist the orthodox Bayesian's claim that a rational agent has a precise credence in every proposition that she can entertain.

Thus defining an agent's credence in terms of her betting quotient has not helped. Intuitively, there is no particular number that is your credence in SARDINES, and similarly there doesn't seem to be any particular number that is your betting quotient for SARDINES — for there is no single number that is *the* number that you would produce were we to elicit your quotient as De Finetti recommends. How then should we respond? In the next section I consider (and set aside) a nihilist position, according to which you have no betting quotient, and no credence in SARDINES. Then in section 3 I set out an alternative theory: that the expression 'credence' is vague.

Nihilism

I start by considering a nihilist position — in order to set this position aside. On this nihilist view, you do not have a credence in SARDINES. We might argue for this as follows:⁴

1. Your credence in SARDINES is the number that you would produce were your betting quotient to be elicited.
2. Thus your credence in SARDINES is the number that you produce in the closet possible worlds in which your betting quotient is elicited.

⁴ This argument rests on the assumption that your credence in SARDINES is identical to your betting quotient in SARDINES, which of course is open to challenge. But it is not clear how the claim that your credence is not identical to your betting quotient can help us here: there are strong intuitive reasons to doubt that you have any particular credence in SARDINES, and the claim that your credence is identical to your betting quotient introduced in the hope that this would fix your credence in SARDINES.

3. But there are a range of equally close worlds in which your betting quotient is elicited, and the number that you produce varies across these worlds.
4. Thus there is no such thing as *the* number that you produce in the closest possible worlds in which your betting quotient is elicited.
5. Thus there is no such thing as your credence in SARDINES.

We could generalize this argument to show that you have no credence in many of the propositions that you entertain. Thus if you have a credence function at all, it does not map each proposition that you entertain onto a number between 0 and 1. Rather, at best it maps *some* of the propositions that you can entertain (such as HEADS) onto numbers between 0 and 1: your credence function does not map propositions like SARDINES to anything whatsoever. Is this a plausible position? And how can we resist the argument for it given above?

I want to start by showing that we can produce an argument paralleling that above to the conclusion that you have no resting heart rate. This seems like a surprising conclusion. Intuitively you *do* have a resting heart rate, even if you are not currently resting. Perhaps you are reading this paper while working out on a treadmill, and so your heart rate is currently elevated, but if so it would still make sense for a trainer to ask what your resting heart rate is — perhaps to check that your current training session is pitched at the right level. Thus intuitively you do have a resting heart rate, even when you are not resting. And this is because your resting heart rate is a dispositional property: it is the rate at which your heart *would* beat *were* you at rest. The problem is that there are a variety of ways to rest, and you could just as easily rest in one way as another. For example, you could rest by lying down in a cool room; or you could rest by sitting up in a warm room. Plausibly there are a range of equally close counterfactual cases where you are resting, and your heart rate will vary across these cases because your heart rate is sensitive to all sorts of factors. There are some guidelines that specify

more precisely what is meant by ‘resting’ (Palatini, 2009), but because these inevitably fall short of exact precision, your heart rate may vary even across close counterfactual cases where the guidelines are met. Thus we can construct the following argument:

1. Your resting heart rate is the rate at which your heart would beat were you at rest.
2. Thus your resting heart rate is the rate at which your heart beats in the closest possible worlds in which you are at rest.
3. But there are a range of equally close worlds at which you are at rest, and the rate at which your heart beats varies across these worlds.
4. Thus there is no such thing as *the* rate at which your heart beats in the closest possible worlds in which you are at rest.
5. Thus there is no such thing as your resting heart rate.

We could also argue along similar lines that you have no height. This is counterintuitive, for surely you do have a height — something in the region of 5’10’, say. This is your height even if you are sitting down, and thus currently measure considerably less than 5’10’ across every dimension. If you are sitting down in the doctor’s office, and she asks for your height, you do not need to check how you are currently oriented to be able to answer. For your height is a dispositional property. It is the distance that there *would* be between the top of your head and the soles of your feet *were* you to stand in a particular way: straight, but not on tiptoe. So your height depends on what this distance is in the closest worlds where you are standing in the relevant way. But the problem is that there will be a range of equally close worlds where you are standing in the relevant way: ones where you stretch your neck out slightly more; ones where you are imperceptibly slouching, and so on. The distance between the top of your head and the soles of your feet will vary across these possible worlds. Measurements of a person’s height vary a surprisingly large amount: the measurement depends on many factors, including

precisely how you are standing, and whether traction is applied as you are measured (Buckler, 1978). Thus there is no such thing as *the* distance that there would be between the top of your head and the soles of your feet *were* you to stand up straight. Thus (following the argument pattern above) we can show that you have no height.⁵

If we accept the conclusions of these arguments, we arrive at a sort of nihilism: you have no height, no resting heart rate, and no credence in many of the propositions that you can entertain. Is this a plausible position? Well, it seems more plausible than Peter Unger's radical position on vagueness (Unger, 1979). Unger has argued that there aren't any tall objects, nor bald objects, nor indeed any 'ordinary things': this is his response to the sorites paradox — a paradox that I discuss in the next section. Unger's position is very counterintuitive, because he denies many claims that we take to be uncontroversially true — such as the claim that a person who is 6'5" is tall. If we deny that you have a heart rate, a height, or a credence in SARDINES, do we similarly deny claims that are uncontroversially true? It is much less obvious. For though at first blush it seems obvious that you have a resting heart rate, say, this intuition is not very robust. Once we start to think about what your resting heart rate is *exactly*, it becomes clear

⁵ With the concept of 'height' there is an added complication. We know that your height is the distance between the top of your head and the soles of your feet in some close possible world(s) in which you are standing in the right sort of way—but we do not know how to choose between the various close possible worlds that seem to meet this criterion. This is the problem described above, which has an obvious analogue for the concepts of 'resting heart rate' and 'betting quotient'. But for the concept of 'height' there is a further problem: we do not even know in any given possible world what the distance is from the top of your head to the soles of your feet. This is because we don't know where your head or feet end—which molecules of dead skin to include—, for the boundaries to your body are vague.

that there is no n such that it is uncontroversial that your resting heart rate is n : thus the intuition that you have a resting heart rate is itself controversial, in contrast to a truly uncontroversial claim, such as the claim that a person who is 6'5" is tall.

Should we then accept that you have no resting heart rate, and no height, and so on? This would leave us with some puzzles. What is happening when the doctor asks for your resting heart rate, and you utter some number, and she takes that number into account in her assessment? If you had no resting heart rate, how could we explain this exchange?⁶ We cannot hope to explain it as make-believe — as we might explain some dialogues involving fictional predicates — for the doctor takes your contribution seriously and uses it to guide her diagnosis. Furthermore, even though there may be no uncontroversial truths of the form *your resting heart rate is n* , there seem to be other uncontroversial truths in the vicinity. For example, it may be uncontroversially true that your resting heart rate is higher than 20bpm, and less than 200 bpm. How can we explain the uncontroversial truth of these claims if you have no resting heart rate?

My suggestion is that we claim that you *do* have a resting heart rate, but that it is a vague matter what this resting heart rate is. The same can be said for your height, and for your credence in SARDINES. In the next section, I consider what it means to say that these predicates are vague — given that they do not fit the usual mould.

⁶ We might try saying that your resting heart rate at any time is just whatever it was most recently measured as — but this is not plausible, because it makes sense to think that your resting heart rate has increased or decreased since it was last measured. And of course we might reasonably form a conjecture about a person's resting heart rate even if we knew that person had never had this measurement taken.

Vagueness

Philosophers usually introduce the phenomenon of vagueness with archetypal one-place predicates, such as ‘... is bald’, ‘... is tall’, ‘... is a small number’, and so on. Vague predicates are taken to have certain features: they typically have both borderline cases and clear cases, appear to lack sharp boundaries, and are susceptible to sorites paradoxes (Keefe, 2000, p. 6).

Can predicates with different forms be vague as well? Rosanna Keefe argues that they can, and that ‘[a] theory of vagueness should have the resources to accommodate all the different types of vague expression’ (Keefe, 2000, p. 14). I claim that predicates such as ‘... at ... has a resting heart rate of...’ are vague expressions — even though they do not seem to have the features typical of vague expressions. Below I work through these features in turn:

1) Vague terms typically have both borderline cases and clear cases

We can use the archetypally vague predicate ‘... is bald’ to illustrate this feature. A person with 0 hairs on his head is clearly bald; a person with 500,000 hairs on his head is clearly not bald; and a person with 5739 hairs on his head may be a borderline case. Thus this predicate has an extension with some objects clearly falling within it, some objects clearly falling outside it, and some borderline cases.

Now let us turn to our predicate ‘... at ... has a resting heart rate (in bpm) of...’ and see whether this predicate also exhibits this feature. This is a three-place predicate: if anything belongs within its extension, it will be ordered triples consisting of an animal, a time, and a number. It is easy enough to find clear cases of ordered triples that do not fall within the extension of this predicate: for example, the ordered triple consisting of myself, now, and the number 20 clearly falls outside the extension of this predicate, because I certainly don’t currently have a resting heart rate of 20bpm. It is also easy enough to find examples of ordered triples that are borderline

cases. For example, the ordered triple consisting of me, now, and the number 75.047 is a borderline case: it is not clearly the case that my current resting heart rate is 75.047bpm, but it is not clearly not the case either. But it seems that we cannot find any examples of ordered triples that clearly fall within the extension of the predicate, for no animal at a time clearly has a resting heart rate of precisely n , for any number n . Thus it seems that this predicate only partially displays this feature of vague predicates.

2) *Vague terms appear to lack sharp boundaries*

Again, let us use ‘... is bald’ to illustrate this feature. Some people fall within the extension of this predicate, and some fall outside. We can then visualize a boundary separating those who are bald from those who are not. But because there are borderline cases, intuitively there is no *sharp* boundary between those who are bald and those who are not. We might instead imagine two boundaries, one separating out the objects that clearly fall within the extension of the predicate, and one separating out the objects that clearly fall outside the extension of the predicate, leaving the objects in between these two boundaries as the borderline cases. But intuitively these boundaries are not sharp either: for there are objects that are not clearly borderline. More generally, archetypal vague predicates seem to lack sharp boundaries altogether.

Does this also hold for predicates such as ‘... at ... has a resting heart rate (in bpm) of...’? Here the whole concept of a boundary is harder to make sense of. Where should we look for a candidate boundary to see whether it is sharp? If we consider all the possible ordered triples of objects, then we will find that a few of these ordered triples are borderline, and the rest fall clearly outside the extension of the predicate. The whole idea of a ‘boundary’ does not seem relevant here. Strictly speaking, then perhaps we can say that this sort of predicate does appear to lack sharp boundaries — simply because it has no boundaries at all. But the way in which

the predicate appears to lack sharp boundaries seems only loosely connected to the way in which ‘bald’ appears to lack sharp boundaries.

3) *Vague terms are susceptible to sorites paradoxes*

We can construct a sorites paradox for the predicate ‘bald’ as follows. We can take some object which clearly falls within the extension of the predicate — such as a person with 0 hairs. Our first premise is then that a person with 0 hairs is bald. Our second premise is a ‘tolerance principle’: for any n , if a person with n hairs is bald, then a person with $n+1$ hairs is also bald. We can use these two premises to argue that a person with 1 hair is bald, and from there that a person with 2 hairs is bald, and so on. Eventually, we reach the conclusion that a person with 500,000 hairs is bald, and this is clearly false, for a person with 500,000 hairs is a clear case of an object that falls outside the extension of the predicate. This argument is a paradox because it has apparently true premises, seems to be valid, but has an apparently false conclusion.

We cannot construct a sorites paradox in the same sort of way for the predicate ‘... at ... has a resting heart rate of...’, for we are missing the item that clearly falls within the extension of the predicate that we need to construct our first premise. Of course sorites paradoxes can be constructed in reverse — so perhaps we could instead start with a premise involving an item that clearly does *not* fall under the extension of the predicate. For example, the first premise could be that I do not currently have a resting heart rate of 20bpm. The obvious difficulty with this strategy is that to construct the paradox in this way we would need some item which clearly falls under the extension of the predicate in order to reach the (apparently false) conclusion. Another difficulty is in giving a persuasive tolerance principle. We might try: for any animal a , time t , and number n , if a at t has a resting heart rate of n bpm, then a at t has a resting heart rate of $n+1$ bpm. But this principle doesn’t even appear true. Thus it seems that predicates like ‘... at ... has a resting heart rate of ...’ are not susceptible to sorites paradoxes.

We have surveyed the three features that vague predicates typically display, and found that none of them are displayed straightforwardly by the predicate ‘... at ... has a resting heart rate of...’. If predicates like ‘... at... has a resting heart rate of...’ are vague, then, they are not archetypal vague predicates. However we might think that we could instead focus on some other closely related predicates that *do* display the typical features of vague predicates. For example, take the two-place predicate ‘... at... has a resting heart rate of at least 100bpm’. First note that this predicate plausibly has both clear cases and borderline cases: some pairs consisting of an animal and a time fall clearly within its extension; some fall clearly outside its extension; and some are borderline. Second, note that this predicate seems to lack sharp boundaries in much the same way that ‘... is bald’ does: some items are borderline, but the set of borderline items is not sharply bounded. Can we pull off a hat-trick, and show that this predicate also gives rise to sorites paradoxes?

Well, we can produce an item that clearly falls within the extension of the predicate, and so can construct the (apparently true) first premise: e.g. the first premise could be the claim that Scrabble the gerbil currently has a resting heart rate of at least 100bpm. We can also produce an item that clearly falls outside the extension of the predicate, and so construct the (apparently false) conclusion: e.g. the conclusion could be the claim that Bradley Wiggins the cyclist currently has a resting heart rate of at least 100bpm. But can we produce a tolerance principle that at least appears to be true? This is tricky. For ‘bald’, we had a variable — the number of hairs — which we could use to construct the tolerance principle. There we could rely on the natural assumption that whether someone is bald depends on the number of hairs that she has. What could play this role in the sorites paradox for ‘... at... has a resting heart rate of at least 100bpm’? What variable does an animal’s resting heart rate necessarily depend upon? There doesn’t seem to be any obvious answer to this question — except of course this trivial answer: the animal’s resting heart rate. But constructing a tolerance principle using this trivial answer

is a non-starter, for (no matter how small we make e) it does not even appear to be true for any n that if an animal at a time with a resting heart rate of n has a resting heart rate of at least 100bpm, then an animal at a time with a resting heart rate of $n-e$ has a resting heart rate of at least 100bpm. Perhaps we might instead imagine arranging all animals-at-a-time in the relevant order, and trying this principle: if any animal-at-a-time in this series has a resting heart rate of at least 100bpm, then the animal-at-a-time immediately to its right also has a resting heart rate of at least 100bpm. But the tricky question here is — what is the relevant order? The idea of course is that the animals-at-a-time should be ranged in order according to their resting heart rate — but how would this be done? The comparative ‘... has a higher resting heart rate than...’ is itself vague.⁷ It is not always clear whether one animal-at-a-time has a higher resting heart rate than another, and so there is no clear ordering of all possible animals-at-a-time according to their resting heart rate. In summary it is not at all obvious how we might construct a sorites paradox for this predicate.

Thus whether we stick with our target predicate ‘... at... has a resting heart rate of...’ or consider closely related predicates such as ‘... at... has a resting heart rate of at least 100bpm’, we find that we are dealing with atypical vague predicates here. In the next section, I consider how we might extend a standard account of vague predicates — supervaluationism — to cover these sorts of predicates. This gives us an account of the (in my view, vague) predicate ‘... at... has a credence in ... of’. I see this account as an alternative to the standard imprecise probabilist’s account, and in section 5 I examine the ways in which these two accounts differ.⁸

⁷ Keefe argues persuasively (Keefe, 2000, p. 12), in disagreement with Neil Cooper (Cooper, 1995), that comparatives can be vague.

⁸ Aidan Lyon (Lyon, 2017) also draws on the vagueness literature to give an account of our doxastic states. Lyon draws inspiration from the degree theorists, who claim that sentences

Supervaluationism

I start by illustrating the supervaluationist's account using the predicate '... is bald'. The supervaluationist claims that there is a range of admissible ways of making the language precise: in other words, there is a range of admissible precise languages, or 'precisifications'. Perhaps under one such precisification, the boundary for '... is bald' lies at 5739 hairs; another might draw the boundary at 5738 hairs — and there will be many other options. For the supervaluationist, a sentence is super-true if and only if it is true under every precisification; it is super-false if and only if it is false under every precisification; and if it is true under some precisifications, and false under others, then it is neither super-true nor super-false.

Thus, for example, the claim that a person with 0 hairs is bald is true under every admissible precisification, and so this claim is super-true. The claim that a person with 500,000 hairs is bald is false under every precisification, so this claim is super-false. The claim that a person with 5739 hairs is bald is true under some precisifications but false under others, so it is neither super-true nor super-false. We can also consider more complex claims, such as the claim that if a person with 5739 hairs is bald, then a person with 5738 hairs is also bald. This is true under

have degrees of truth between 0 and 1. Though Lyon's account is inspired by the degree theorists, his account is really a new (and very interesting) extension of the machinery that the degree theorists use to give their account of vagueness. In contrast, I am suggesting that no new account is needed: a theory of vagueness (and in my paper I focus on supervaluationism rather than the degree theory) *already* has the resources to explain the intuitions that seem to drive us away from the orthodox Bayesian account. Many thanks to Dan Lassiter for bringing Lyon's paper to my attention.

every precisification, and so counts as super-true. Similarly, this disjunction comes out as super-true, even though neither disjunct is super-true: either a person with 5739 hairs is bald, or she isn't. More generally, all instances of the law of the excluded middle — and indeed all the laws of classical logic — come out super-true. Furthermore, even though there is no number n such that it is super-true that a person with n hairs is bald but a person of $n+1$ hairs is not, nevertheless the claim that there is some such n is super-true — for it is true under each precisification. This gives the supervaluationist a solution to the sorites paradox: the tolerance principle is false under every precisification, and therefore super-false.

Can we extend this account to give a supervaluationist account of our predicate '... at ... has a resting heart rate (in bpm) of...'? A natural move here is to say that under each precisification of the language, this expression denotes a relation, and the relation consists of ordered triples — each containing an animal, a time, and a number. For each animal and each time, there will be only one number that appears in a triple with that animal and time in the relation: we can then say that the animal and time pair are mapped onto, or assigned, that number. Thus under one admissible precisification of the language, the expression '... at... has a resting heart rate of...' might assign me, right now, the number 75.38: under this precisification, it is true that my resting heart rate is 75.38bpm. Under a different precisification, the expression will pick out a different relation, which perhaps assigns a different number to me.

Let us say then that under each admissible precisification of the language, the expression '... at ... has a resting heart rate of...' refers to a relation that assigns a single number to every animal (with a heart) at every time. Whether a sentence is super-true, super-false, or neither depends on whether the sentence is true under all, none, or some of these precisifications of the language. Thus the sentence 'my resting heart rate now is at least 20bpm' will come out super-true, because it will be true under every precisification; the sentence 'my resting heart rate now

is at least 120bpm' will come out super-false because it is false under every precisification; and a sentence like 'my resting heart rate now is exactly 75.384' will come out neither super-true nor super-false, because it will be true under some precisifications but false under others.

Here it might be objected that there is an important difference between the supervenient account of the vague term '... is bald', and the supervenient treatment that I offer of '... at ... has a resting heart rate of...'.⁹ The supervenient claims that there are various admissible precisifications of '... is bald', and (the objector might say) we can give a straightforward intuitive definition for each such precisification. For example, take the precisification according to which 'bald' means bald₅₇₃₉: we can define this precisification straightforwardly by saying that a person is bald₅₇₃₉ iff she has 5739 or fewer hairs on her head. In contrast (the objector might continue), the same does not hold for my supervenient treatment of '... at... has a resting heart rate of ...'. I claim that there are ways of making this expression precise, but what are these precisifications? How might we define them? The first part of my response is to clarify that it is not just one-place predicates like '... is bald' that can be vague. Predicates with more than one place can be vague too, and given a supervenient treatment. To use Keefe's example, each admissible precisification of the predicate '... is a friend of...' has an extension that is a set of ordered pairs (Keefe, 2000, p. 159), and similarly, I claim that each admissible precisification of the predicate '... at... has a resting heart rate of ...' has an extension that is a set of ordered triples (each triple consisting of a person, a time, and a number). The second part of my response is to clarify that it is not part of the supervenient story that there must be a straightforward and intuitive way to characterize each precisification: 'all that is needed for the truth-conditions is the range of precise extensions themselves' (Keefe, 2000, p. 159). If the supervenient did require a straightforward and intuitive characterization of each

⁹ Thanks to the audience at Stirling Philosophy Department for pushing this objection.

precisification, then the supervaluationist could not apply her account to vague terms such as ‘... is nice’ — or indeed ‘... is a friend of...’. Thus, in response to the objection, I claim that my supervaluationist treatment of the predicate ‘... at ... has a resting heart rate of ...’ is not importantly different from the supervaluationist treatment of many other expressions.

Having seen how we can give a supervaluationist account of the predicate ‘... at ... has a resting heart rate of...’, let us now give a similar supervaluationist account of the predicate ‘... at... has a credence in ... of ...’. Under each precisification of the language, this expression will denote some relation. This relation will consist of ordered 4-tuples, each containing a person, a time, a proposition and a number. For any person, time, and proposition, there will be only one number that appears with them in a 4-tuple in the relation, and so we can say that the combination of person, time, and proposition are mapped onto (or assigned) a particular number. Thus for example, under one precisification of the language, the expression will refer to a relation that maps you, now, together with the proposition SARDINES, onto the number 0.342: under this precisification, it is true that your credence in SARDINES right now is 0.352. Under another precisification, it is true that your credence right now in SARDINES is 0.353. Once again, we say that a sentence is super-true iff it is true under every precisification; super-false if false under every precisification, and neither super-true nor super-false if true under some but not all precisifications.¹⁰

¹⁰ It follows on this account that for any pair of propositions P and Q, and any rational agent, the following disjunction will be super-true (because true under all precisifications): either the agent has a higher credence in P than in Q, or the agent has a higher credence in Q than in P, or the agent has the same credence in P as in Q. As Richard Dietz pointed out to me, this is in conflict with the views of some theorists (Keynes, 1921)—though we can soften the conflict

This, then, is one way to defend orthodox Bayesianism: claim that the expression ‘credence’ is vague and so covered by the supervaluationist’s theory of vagueness.¹¹ The orthodox Bayesian’s claim — that every rational agent has a precise credence in every proposition entertained — then emerges as super-true (because it is true under every admissible precisification), but it does not have the unwelcome consequences that we might expect. For example, on this view, orthodox Bayesianism is compatible with there being no particular n such that it is super-true that your credence in SARDINES is n . This, then, is one way that the orthodox Bayesian could reconcile her account with our intuitions. In the next section, I turn to an alternative: the imprecise probabilism account. This influential account is also sometimes

by adding that on the supervaluationist account it may be that neither of the disjuncts are super-true.

¹¹ I claim that the expression ‘credence’ is vague, but I do not claim that it is vague along every dimension. For example, take the predicate ‘... is a credence function’. At least on some views, this predicate is precise, for everything either is a credence function (i.e. a function—perhaps conforming to certain axioms, depending on your preferred definition of ‘credence function—from propositions to numbers between 0 and 1) or it isn’t. My claim is rather that it can be vague whether a particular agent at a time has a particular credence function. Furthermore, the idea is not simply that it might be vague whether a particular agent has *any* credence function (perhaps because it is vague whether the agent is capable of being in a belief state), or that it might be vague what relation the agent stands in to the function (i.e. whether it is an epistemic relation rather than some other closely related relation). The idea is that it can be vague *which* credence function is the credence function of a particular agent at a specific time. Thanks to the audience at the philosophy department in Sterling for pushing me to clarify these points.

referred to as a type of supervaluationism (Hajek, 2003, pp. 277–8; van Fraassen, Vague Expectation Value Loss, 2006, p. 483), but it is an account inspired by supervaluationism, rather than a mere application of the standard theory. I describe this imprecise probabilism account in the next section, and then compare it with the supervaluationist account that I have given above.

Imprecise Probabilism

On the orthodox Bayesian view, a rational agent's credal state can be represented by a single credence function. Many theorists, including (Jeffrey, 1983; van Fraassen, 1990; Joyce, 2010), deny this and claim instead that a rational agent's credal state can be represented by a *set* of precise credence functions. The credence functions in the set will be such that whatever holds of the agent's credal state holds for every credence function in this set: or as Joyce puts it, '[f]acts about the person's opinions correspond to properties common to all the credence functions in her credal set' (Joyce, 2010, p. 287) Thus, for example, suppose that your credence in SARDINES is greater than 0.2: in that case every credence function in the set representing your credal state must assign a value of at least 0.2 to SARDINES. Suppose also that your credence in SARDINES is less than 0.8: in that case, every credence function in the set representing your credal state must assign a value of less than 0.8 to SARDINES. It is natural to think of your credence in SARDINES as a range — given by the range of numbers that the credence functions in the set representing your credal state assign to SARDINES. We can also consider more complex claims about your credal state. For example, it may be that your credence in SARDINES is greater than your credence in EXTREME SARDINES, where EXTREME SARDINES is the claim that my neighbour has more than ten tins of sardines in her cupboard. This claim about your credal state will obtain iff every precise credence function

in the set representing your credal state assigns a higher number to SARDINES than to EXTREME SARDINES.

Having given this brief overview of imprecise probabilism, I turn to a comparison with the supervaluationist account given in the previous section. The fundamental difference between the accounts is that for the supervaluationist, the orthodox Bayesian theory is true but expressed in vague language; whereas for the imprecise probabilist, the orthodox Bayesian theory is false and should be replaced with imprecise probabilism, on which we have alternative models that represent epistemic states precisely. In the following three sections I compare the accounts by looking at three different features: I consider how each account handles complex claims; then I consider how a decision theory can be constructed on each account; and finally I look at how each account can handle higher-order vagueness (or its analogue).

Complex Claims

Here is an important difference between the two accounts. The imprecise probabilist works with sets of precise credence functions, and each such set of credence functions is designed simply to represent the credal state of a particular agent at a particular time.¹² In contrast, the

¹² On some versions of imprecise probabilism, we represent an agent's mental state with a set of *pairs*, each pair consisting of a credence function and a utility function. A set of such pairs represents more than just the agent's credence function—but there is still an enormous contrast with the supervaluationist's set of precise languages, which can represent far more than the mental state of a single agent at a time.

supervaluationist works with a set of precise languages, and these languages can be used to say all sorts of things about all sorts of topics.

For the supervaluationist, each precisification of our language contains a precisification of the expression 'credence'. We can interpret a precisification of 'credence' (as we did earlier) as a 4-place predicate that we can see as assigning numbers to various combinations of individual, time, and proposition. For example, a given precisification of 'credence' might assign the number 0.342 to the combination of me, the present time, and the proposition SARDINES, and also assign the number 0.341 to the combination of you, the present time, and the proposition SARDINES. To put it another way, under this precisification of the language, it is true that my credence right now in SARDINES is 0.342, and also true that your credence right now in SARDINES is 0.341. Thus these claims about both your and my credal states each get a truth-value under this single precisification of the language. Similarly, claims about the relation between your credal state and mine get a truth-value under each precisification of the language. For example, take the claim that my credence (now) in SARDINES is greater than yours (now). This claim will have a truth-value under each precise language: for example, on the precisification just mentioned, under which it is true that my credence right now in SARDINES is 0.342, and true that your credence right now in SARDINES is 0.341, it will be true that my credence in SARDINES is greater than yours. In this way claims relating multiple credal states will have a truth value under any given precisification. And on the supervaluationist account, we can go on to say that a claim relating multiple credence functions will be super-true if and only if it is true under every precisification, super-false if and only if it is false under every precisification, and otherwise neither super-true nor super-false. Thus, on the supervaluationist account, there is no difficulty in explaining how truth-values are assigned to claims relating different credal states.

What about the imprecise probabilist's account? Well, here we do not have a set of precise languages, but rather sets of precise credence functions. Your current credence function is represented by one such set, and my current credence function is represented by some other such set. It makes no sense to talk about what my credence is under one of the precise credence functions that represent your credal state, or vice versa. The truth-value of claims about your credal depend on the set of credence functions that represent your credal state; the truth-value of claims about my credal state depend on the set of credence functions that represent my credal state. This leaves the imprecise probabilist with a question: what settles whether a claim holds when that claim is about the relationship between our credal states?

The imprecise probabilists have not, as far as I know, generally agreed on an answer to this question, but there are at least two sorts of answers that the imprecise probabilist could give:

- (i) The imprecise probabilist could say that whether my credence in SARDINES is greater than yours depends simply on our credal ranges in these two propositions. For example, the condition could be that the highest value in my credal range must be greater than the highest value in your credal range. Alternatively — or additionally — it might be required that the lowest value in my credal range must be greater than the lowest value in your credal range. An alternative and more demanding requirement could be that the lowest value in my credal range must be greater than the highest value in your credal range.¹³

¹³ This leaves open the question of when we should say that my credence in SARDINES is equal to yours—and this is a further issue for the imprecise probabilist to resolve. One option here is to say that all that is required is that my credence in SARDINES must be neither lower nor higher than yours—which (on the ‘demanding requirement’ in the text above) would mean

- (ii) The imprecise probabilist could say that whether my credence in SARDINES is greater than your credence in SARDINES depends on whether a particular sort of mapping relation holds between the credence functions in each of our sets. To understand this idea, imagine trying to one-to-one map each credence function in your set to some corresponding credence function in my set – the only criterion being that for each credence function in your set, the number assigned to SARDINES must be less than the number assigned to SARDINES by the corresponding credence function in my set. If some such mapping is possible, then my credence in SARDINES is higher than yours. We can similarly define what it is for my credence in SARDINES to be lower than yours, and it is natural to extend the account to define what it is for my credence in SARDINES to be equal to yours — which would be namely, for there to be a function that maps each credence

just that there needs to be some overlap between your credal range and mine. Perhaps, though, this makes equality implausibly easy to come by. In contrast, we might rule that my credence in SARDINES is equal to yours only if our credal ranges are identical, and this may make equality implausibly hard to come by. There is an interesting parallel here with Joshua Gert's account of value (Gert, 2004), and I am grateful to Richard Dietz for pointing me to this account.

function in your set to some corresponding credence function in my set that assigns the very same number to SARDINES.^{14,15}

The question that we have put to the imprecise probabilist may not seem particularly pressing. Perhaps the whole idea of my credence being greater than yours in SARDINES is rather obscure, given that our credences in this claim are intuitively imprecise. What would it mean for my credence function to be higher than yours in this context? But there is a related question that does seem more pressing, and this concerns comparisons of the same person's credal state across time. For example, suppose that you come to learn NOT EXTREME SARDINES — which is the claim that my neighbour does *not* have more than ten tins of sardines in her cupboard. Intuitively, on learning this your credence in SARDINES ought to decrease, for one of the ways that SARDINES could be true has just been eliminated. But what makes it the case that your credence in SARDINES is higher before gaining this evidence than it is after? It is not immediately obvious how the imprecise probabilist should answer this question, but we may be able to construct an answer by thinking about the imprecise probabilist's approach to rational updating, which I turn to now.

¹⁴ It follows on this account that for my credence in SARDINES to be greater than yours (or lower than yours or equal to yours) the sets of credence functions that represent each of our credal states need to agree in cardinality. This seems like a very strong constraint, and this may be an objection for an imprecise probabilist who takes this option. Thanks to Richard Dietz for pointing this out.

¹⁵ I suspect that—given certain natural constraints—this second option will collapse into some version of the first option. We will see the point of introducing the mapping relation below when discussing conditionalization.

On the orthodox Bayesian view, a rational agent updates by conditionalizing on any new evidence that she encounters. Thus if a rational agent encounters (only) some new piece of evidence E , then her new credence function will be identical to her old credence function conditionalized on E . What does the imprecise probabilist require of a rational agent who encounters only E ? For the imprecise probabilist, the credal states of the agent at the earlier and later times will each be represented by a set of credence functions (rather than a single credence function). What should the relationship between these two sets be, if the agent has updated on new evidence E , as rationality demands? The answer the imprecise probabilists gives¹⁶ is that there should be a particular sort of one-to-one mapping relation between the two sets: we should be able to map each function from the earlier set onto some function from the later set, such that the function from the later set is identical to the corresponding function from the earlier set conditionalized on E .¹⁷

In describing updating on the imprecise probabilist's account, it is tempting to think of the individual credence functions enduring between the earlier and the later time — as though when the agent encounters E , what happens is that each individual precise credence function is updated. This impression is reinforced by the common use of a (vivid and helpful) metaphor, in which we think of the agent's credal state as a set of 'avatars' (Bradley, 2009) or 'committee members' (Joyce, 2010), each with a precise credence function that gets updated as the agent whose credal state they represent gains evidence. Here it is natural to think of each avatar enduring a change in its credence function — but this was always intended just as a metaphor.

¹⁶ This is at least the answer usually given—but other updating rules are possible. See (Joyce, 2010, pp. 292–3) for a discussion of an alternative.

¹⁷ 'If a person in credal state C learns that some event D obtains (and nothing else), then her post-learning state will be $C_D = \{c(\bullet|D) = c(X) \cdot [c(D|X) / c(D)] : c \in C\}$ ' (Joyce, 2010, p. 287).

If we focus instead on the sets of credence functions themselves, then it is clear that it does not make sense to think of a single credence function changing its assignment of values over time: functions that assign different values are simply different functions. Thus, in a case of rational updating, there is no identity relation that connects each credence from the earlier set with some credence in the later set. On the imprecise probabilism view, then, the relationship that must hold for an agent to have updated rationally in response to new evidence must be as I described above: the condition is that there must be some one-to-one mapping relation between the two sets such that each function from the earlier set is mapped onto a function from the later set, which assigns every proposition whatever number the function from the earlier set assigned to that proposition conditional on E.

Given that this mapping relation is what is needed for an agent to count as having updated rationally, it is natural to also use mapping relations to explain relations between different credal states more generally. Thus we can say that for your credence in SARDINES to decrease after learning EXTREME SARDINES is for a certain mapping relation to hold between the credence functions in the set that represent your earlier credence function and those in the set that represent your later credence function: there must be some one-to-one mapping of each credence function from the earlier set onto some credence function in the later set, such that the credence function from the earlier set assigns a higher value to EXTREME SARDINES than the credence function in the later set. And we can tell a similar story about what must be the case for your credence in SARDINES to be lower than mine—along the lines of (ii) above. Thus it is possible for the imprecise probabilist to give an account of these sorts of complex statements, but the account does not follow automatically from the core claims of imprecise probabilism.¹⁸

¹⁸ See footnotes 13 and 14 for some of the difficulties in giving such accounts.

For the supervaluationist, the story is pretty straightforward. Under any given precisification, it will be either true or false that your credence function in SARDINES is lower than mine; that my credence in SARDINES has decreased since learning EXTREME SARDINES; or that on learning some evidence I have conditionalized as rationality demands. Whether these claims are super-true, super-false, or neither simply depends on whether the claims are true under all, none, or some of the precise languages. Thus whether we go for a supervaluationist account, or imprecise probabilism, we can make sense of claims that relate one credal state to another—but while this follows automatically for the supervaluationist, the details need to be bolted on to imprecise probabilism.

Decision Theory

As we have seen, the supervaluationist works with a set of precise languages, which can be used to say things about all sorts of topics; in contrast, the imprecise probabilist works with sets of credence functions, and each set is used just to represent a single agent's credal state at a time. We have seen how this leads to a difference in the way that the two accounts handle claims about relations between the credences of different people, or about the relations between the same person's credence at two different times. It also leads to a difference in the way that decision theory works on the two accounts.

The imprecise probabilists have put forward many different decision theories, but here I just briefly describe two sample theories:¹⁹

¹⁹ I focus on normative theories, but there are also attempts to produce descriptive theories.

1. *Caprice* (Weatherson, 2008): on this theory, an agent is rationally required to maximise expected utility relative to at least *one* credence function in the set that represents her credal state. Thus, for example, suppose that I offer you a bet on SARDINES, whereby you pay out £0.35 and get £1 back iff SARDINES. Let us assume that you value only money, and value that linearly. Then you are permitted to accept this bet iff there is a credence function within the set that represents your credal set that assigns SARDINES a value of at least 0.35. Similarly, you are permitted to reject this bet iff there is a credence function within the set that assigns SARDINES a value of no more than 0.35. There may be many bets that you are rationally permitted to either accept or reject (at your caprice).
2. *Maximin* (Gärdenfors & Sahlin, 1982) (Gilboa & Schmeidler, 1989): On this theory, an agent is rationally required to maximise minimum expected utility. To understand what this means, let us again suppose that I offer you a bet on SARDINES, whereby you pay out £0.35 and get £1 back iff SARDINES is true. You have the option of either accepting the bet or rejecting it. We can start by calculating the minimum expected utility of accepting the bet. To do this, we calculate the expected utility of accepting the bet relative to each credence function in your credal set: the lowest expected utility that we get from this process is the minimum expected utility of accepting the bet. We then calculate the minimum expected utility of rejecting the bet in a similar way. Thus the actions available to you (accepting the bet, rejecting the bet) each have a corresponding minimum expected utility, and if one has a higher minimum expected utility than the other, then you are required to choose that action. More generally, in any choice situation you are rationally required to select from amongst those actions with the maximum minimum expected utility.

Various problems face both of these decision rules, and I discuss some of these problems elsewhere (Mahtani, 2016). But the point that I want to note here is that both of these decision rules refer to the *set* of credence functions that represent the agent's credal state.

For the supervenientist, in contrast, the most natural thought is that a given decision rule will be stated in the object language, and will make no reference to the set of admissible precisifications. For example, the supervenientist might endorse the simple decision rule MEU, according to which a rational agent always chooses an action that maximises expected utility. On the supervenientist account, different precisifications of the language precisify 'credence' differently, and so the actions that can truly be said to maximise expected utility, and therefore be rational, can vary across precisifications. For example, suppose that I offer you a bet whereby you pay out £0.35 and get £1 back iff SARDINES obtains. We assume as before that you value only money, and value it linearly. Should you then accept or reject this bet? Which action has higher expected utility? Well, the expected utility of the actions depend on your credence in SARDINES, and 'credence' is vague and so can be made precise in different ways. Under some precisification, perhaps it is true that your credence in SARDINES is 0.351, and so the expected utility of accepting the bet is greater than the expected utility of rejecting it, and so it would be rational to accept the bet. Under another precisification, perhaps it is true that your credence in SARDINES is 0.349, and so the expected utility of accepting the bet is less than the expected utility of rejecting it, and so it would not be rational for you to accept the bet. Thus it may be true under one precisification that an action is rational, but false under another. It is super-true that an action is rational only if the claim that it is rational is true under every precisification; it is super-true that an action is not rational only if the claim that it is not rational is true under every precisification. If the claim that an action is rational is true under some but not all precisifications of the language, then it is neither super-true that the action is rational nor super-true that the action is not rational: we can then say that it is

indeterminate whether the action is rational. Let us call this the supervaluationist's simple decision theory.²⁰

Here supervaluationism seems to give us a simpler decision theory than imprecise probabilism.²¹ The imprecise probabilist needs to construct some new sort of decision theory to accommodate the fact that an agent's credal state is represented by a set of credence functions rather than a single credence function. In contrast, the supervaluationist can stick with decision rules developed by the orthodox Bayesian, such as MEU. But does the supervaluationist's simple decision theory have the right implications?

One problem for the supervaluationist's simple theory is that in some cases an agent may be faced with a decision problem in which there is no action such that it is super-true that that action is rational. To see this, consider again the decision problem above in which you must choose between accepting and rejecting the bet for which you pay out £0.35 and get £1 back iff SARDINES is true. Is it super-true that it is rational to accept the bet? Well, if under some admissible precisification it is true that your credence in SARDINES is 0.349, then under this precisification it is false that it would be rational to accept the bet, and so it is not super-true that it would be rational to accept the bet. Is it then super-true that it is rational to reject the bet? Well, if under some admissible precisification it is true that your credence in SARDINES is 0.351, then under this precisification it is false that it would be rational to reject the bet, and so it is not super-true that it would be rational to reject the bet. Thus it may be that neither the claim that it is rational to accept the bet, nor the claim that it is rational to reject the bet, come

²⁰ Robert Williams briefly explores a related idea in (Williams, 2014, pp. 25–26).

²¹ Of course, much depends on whether this simpler decision theory is plausible—a question that I explore in the current section.

out as super-true. Whichever action you perform, then, it will not be super-true that your action is rational. This looks like a problem: it seems that in any given decision problem there should be at least one action such that it is super-true that choosing this action is rational.

Another problem for the supervaluationist's simple decision theory is the phenomenon of 'ambiguity aversion'. This can be demonstrated with the 'Ellsberg Paradox' (Ellsberg, 1961). Suppose that you know that an urn contains 30 red balls, and 60 balls that are either black or yellow: you do not know the ratio of black to yellow balls. You are about to draw out a ball at random, and are given a choice between these two options: (a) you receive £100 if you draw a red ball; (b) you receive £100 if you draw a black ball. Now suppose instead that you are given the choice between these two options: (c) you receive £100 if you draw a red or yellow ball; and (d) you receive £100 if you draw a black or yellow ball. Many apparently rational people prefer (a) to (b), but prefer (d) to (c). This is the phenomenon of 'ambiguity aversion': as the imprecise probabilist would put it, a rational agent may prefer an action where every credence function in the set that represents her credal state assigns this action the same expected utility, over an action where the expected utilities assigned differ—all else being equal. Some of the decision rules put forward by the imprecise probabilist are designed to explain this phenomenon: the rule maximin outlined above is one such rule. But can the phenomenon be accommodated on the supervaluationist's account?

On the simple supervaluationist's decision theory that we are exploring here, under each precisification of the language the agent has some particular credence in (RED) the ball drawn being red, some particular credence in (BLACK) its being black, and some particular credence in (YELLOW) its being yellow. Presumably on some precisifications of the language, the agent has a higher credence in RED than BLACK, while on others the agent has a higher credence in BLACK than RED. Thus under some precisifications of the language (those on which the

agent has a higher credence in RED than BLACK) the agent prefers (a) to (b) and (c) to (d), while under other precisifications of the language (those on which the agent has a higher credence in BLACK than RED), the agent prefers (b) to (a) and (d) to (c). But this leaves some claims about the agent's preferences—e.g. the claim that the agent prefers (a) to (b)—indeterminate, i.e. neither super-true nor super-false. And it leaves the complex claim—that the agent prefers (a) to (b) and (d) to (c)—super-false. And this just doesn't seem right: many rational people report an apparently definite preference for (a) over (b) and (d) over (c).

To handle these two problems, the supervaluationist can move away from her simple decision theory to a theory that involves the terms 'super-true'/'super-false', and 'admissible precisifications'. These terms are part of the supervaluationist's meta-language — that is, the language that the supervaluationist uses to express her theory. Thus, if the decision theory uses these terms, then the decision theory will also be in the meta-language. This may seem strange, because surely the meta-language is designed to express supervaluationism rather than to express a decision theory. However this is not really such a big departure for supervaluationism, for supervaluationists typically ascend to the meta-language to justify their position, and this can involve showing how supervaluationism can successfully predict and explain verbal behaviour. For example, the supervaluationist can explain why, when confronted with a borderline red object and asked whether it is red, people will often refuse to assert either that the object is red or that object is not red, perhaps saying instead that there is no fact of the matter: the supervaluationist can explain this by claiming that this is how people do (and perhaps should) respond in cases where neither a statement nor its negation are super-true. Thus the supervaluationist typically ascends to the meta-language to explain and predict verbal behaviour, and so it is no great jump for the supervaluationist to ascend to the meta-

language in order to explain and predict behaviour in general—i.e. in order to give a decision theory.²²

By ascending to the meta-language, we can express some alternative decision theories that the supervaluationist could adopt. For example, the supervaluationist could claim that an action is rationally permissible if and only if that there is at least precisification under which it is true that the action maximises expected utility. This gives us a decision rule parallel to the imprecise probabilist's rule 'caprice'. Alternatively, the supervaluationist could claim that an action is rationally permissible if and only if it has the highest minimum expected utility—where an action's minimum expected utility is the lowest expected utility that it is assigned under any precisification. This gives us a decision rule parallel to the imprecise probabilist's rule 'maximin', which can explain and predict ambiguity aversion.

In summary, then, the supervaluationist can offer a range of different decision theories. There is the supervaluationist's simple decision theory, which does not have a parallel in the imprecise probabilist's account. In addition, by ascending to the metalanguage the supervaluationist can offer a range of decision rules, which parallel those offered by the

²² An alternative for the supervaluationist may be to use the term 'determinately' (or 'definitely') to state her decision theory. To define this term: for any claim p , if the claim p is super-true, then the claim 'determinately p ' is true under every precisification; if the claim p is not super-true, then the claim 'determinately p ' is false under every precisification. The idea is that these terms are part of the object language, and allow us to express what we would otherwise have to express in the meta-language. It may be (though I have not shown this) that whatever decision rule can be expressed in the meta-language can equally well be expressed in the object language, supplemented with 'determinately'.

imprecise probabilist. Thus the supervaluationist has resources at least as good as imprecise probabilism to offer a plausible decision theory.

Higher-Order Vagueness

Supervaluationism faces the problem of higher-order vagueness. Recall that intuitively there is no sharp boundary between the possible objects that are bald and those that are not bald: this is what motivates supervaluationists to construct their account. But the supervaluationist's account seems to avoid positing one sharp boundary, only to posit two more: one marking off the objects that it is super-true to describe as bald from the borderline cases, and another marking off the objects that it is super-false to describe as bald from the borderline cases. Intuitively there are no sharp boundaries here either—and in fact 'bald' does not draw any sharp boundaries whatsoever. How can the supervaluationist handle this?

A convincing response to this problem is to claim that the metalanguage — the language in which the theory of supervaluationism is expressed—is itself vague. In particular, the expression 'admissible precisification' is vague: there is no sharp boundary between the precisifications that are admissible and those that are not, and so it follows that there is no sharp boundary between those claims that are super-true/-false and those that are not. The supervaluationist can then give the same account of the vagueness in the metalanguage as she gave of the vagueness in the object language: there are many admissible precisifications of the metalanguage, and whether a statement in the metalanguage is super-true will depend on

whether the sentence is true under each precisification of the metalanguage.²³ Of course we will then want to say that the meta-meta-language is also vague, and so on, and we will end up with an infinite hierarchy of vague metalanguages—but this is not obviously problematic (Keefe, 2000, pp. 202–213). Thus the supervaluationist has at least one promising response to the problem of higher-order vagueness.

The imprecise probabilist faces a problem analogous to higher-order vagueness (Maher, 2006; Kaplan, 2010; Rinard, forthcoming). For the imprecise probabilist, the problem can be put as follows. Intuitively you do not have a precise credence in SARDINES, for you do not know what your credence is in this claim, and there doesn't seem to be anything that makes it the case that your credence is one value rather than another. The imprecise probabilist handles this problem by claiming that your credal state is represented by a set of credence functions, and these will assign between them a range of values to SARDINES. But what is this range, exactly? What is the lowest number in this range, and what is the highest? The problem with the orthodox Bayesian theory seems to resurface in a different form. For you don't know what numbers form the upper and lower bound of your credal range. And there doesn't seem to be anything that could make it the case that these upper and lower bounds are any particular numbers. How can the imprecise probabilist handle this problem?

One option for the imprecise probabilist is to claim that the account of imprecise probabilism is given in a vague language, and then to use something like supervaluationism to give an account of this vagueness. Thus the imprecise probabilist could claim perhaps that it is vague

²³ This sentence—which gives an account of the vagueness of the meta-language—is itself in the meta-meta-language, and the expression 'super-true' cannot be assumed to have the very same meaning in the meta- and meta-meta-languages.

whether a particular credence function belongs in the set that represents an agent's credal state. This approach is not very appealing, however, because it makes the imprecise probabilist's position quite complicated, with different sorts of accounts needed to handle the original problem and the 'higher-order' problem. The imprecise probabilist would have to claim both that the standard Bayesian model should be replaced by an alternative model, and then that the account describing the alternative model is vague and should be given a supervenientist treatment. It would be simpler to give a thorough-going supervenientist account, as I have done, and drop the claim that the standard Bayesian model needs to be replaced.

As an alternative, the imprecise probabilist might try to iterate her original account at higher-levels, just as the supervenientist gives a supervenientist treatment of vagueness at every level. To work out how this might look, recall that the imprecise probabilist recommends that we reject the standard Bayesian account on which an agent's epistemic state is modelled with a single credence function, in favour of an account on which an agent's epistemic state is modelled with a set of credence functions. Then, in response to a problem analogous to higher-order vagueness, the imprecise probabilist might claim that this account, which models an agent's credal state with a set of credence functions, should in turn be replaced by an account that models an agent's credal state with a set of sets of credence functions. This account will then be replaced by a still more complex account—and so on. But this response lacks the appeal of the supervenientist's response to higher-order vagueness. The imprecise probabilist originally *rejected* the standard Bayesian account in favour of her alternative. Thus, presumably each iteration of the imprecise probabilist's account should similarly be rejected in favour of the next more complex account—meaning that every account in this series should

be rejected.²⁴ Perhaps we should accept an account on which an agent's credal state is modelled by an infinitely long series of sets of credence functions—but what use could we have for such an account? The supervaluationist's response to higher-order vagueness does not suffer from the same problem. The original supervaluationist account of vagueness in the object language is perfectly in order as it is, and the phenomenon of higher-order vagueness gives us no reason to reject it. The account is expressed in vague language, but that is no reason to reject it: the supervaluationist can maintain that it is the correct account of vagueness in the object language. It is only if we want an account of vagueness in the meta-language that we need to ascend into a meta-meta-language to give this account.

The supervaluationist then has a plausible response to the problem of higher-order vagueness. The imprecise probabilist faces an analogue of this problem, and I have looked at some ways in which the imprecise probabilist might respond—but there are problems with each option considered.

Conclusion

Orthodox Bayesianism—which rules that every rational agent has a precise credence in every proposition that she entertains—is counterintuitive. Imprecise probabilism has proved a popular refinement of orthodox Bayesian, but I have argued that there is a viable alternative. The alternative involves recognizing that the expression 'credence' is vague—just as 'resting heart rate' is vague. We can then use supervaluationism, more or less in its standard form, to

²⁴ This criticism is related to the criticism that Keefe levels against a version of the degree theory of truth (Keefe, 2000, pp. 117–120).

give an account of these expressions. Using this account, I have offered a simple way of handling sentences relating the credences of different people, or of the same person at different times; I have also suggested some options for constructing a decision theory; and I have shown how we can make use of a convincing supervenient approach to higher-order vagueness.

Part of the appeal of replacing imprecise probabilism with a supervenient account is that the supervenient account comes for free. Vagueness is a pervasive phenomenon, and we need to give an account of it. Supervenientism is arguably the best account of vagueness that we have, and so every vague term should be given a supervenient treatment. That the expression ‘credence’ is vague is hardly controversial: pretty much all the terms in our language are vague, and there is no particular reason why ‘credence’ should be an exception. All of this implies that—for reasons unconnected to epistemology—we should give a supervenient treatment of the expression ‘credence’. I have argued that once we give a supervenient treatment of this term, we get everything that imprecise probabilism offers and more, and so imprecise probabilism should be abandoned as unnecessary.

References

Bradley, R. (Forthcoming). *Decision Theory with a Human Face*.

Bradley, R. (Forthcoming). Ellsberg's Paradox and the Value of Chances. *Economics and Philosophy*.

Bradley, R. (2009). Revising Incomplete Attitudes. *Synthese*, 171 (2), 235-256.

Buckler, J. M. (1978, Sep). Variations in Height Throughout the Day. *Archives of Diseases in Childhood*, 53 (9), p. 762.

Cooper, N. (1995). Paradox Lost: Understanding Vague Predicates. *International Journal of Philosophical Studies*, 3, 244–69.

- De Finetti, B. (1964). Foresight: Its Logical Laws, Its Subjective Sources. In H. E. Kyburg, & H. E. Smokler, *Studies in Subjective Probability* (pp. 93–158). New York: Wiley.
- Ellsberg, D. (1961). Risk, Ambiguity, and the Savage Axioms. *The Quarterly Journal of Economics*, 75 (4), 643–669.
- Gärdenfors, P., & Sahlin, N.-E. (1982). Unreliable Probabilities, Risk Taking and Decision Making. *Synthese*, 53, 361–386.
- Gert, J. (2004). Value and Parity. *Ethics*, 114 (3), 492–510.
- Gilboa, I., & Schmeidler, D. (1989). Maxmin Expected Utility with Non-Unique Prior. *Journal of Mathematical Economics*, 18, 141–153.
- Hajek, A. (2003). What Conditional Probability Could Not Be. *Synthese*, 137 (3), 273–323.
- Horgan, T. (Forthcoming). Troubles for Formal Epistemology. *Res Philosophica*.
- Jeffrey, R. (1983). Bayesianism with a Human Face. In J. Earman, *Minnesota Studies in the Philosophy of Science, Volume 10: Testing Scientific Theories* (pp. 133–156). Minneapolis: University of Minnesota Press.
- Joyce, J. (2010). A Defense of Imprecise Credences in Inference and Decision Making. *Philosophical Perspectives*, 24, 281–323.
- Kaplan, M. (2010). In Defense of Modest Probabilism. *Synthese*, 176 (1), 41–55.
- Keefe, R. (2000). *Theories of Vagueness*. Cambridge: Cambridge University Press.
- Keynes, J. M. (1921). *A Treatise on Probability*. London: Macmillan & Co., Ltd.
- Lyon, A. (2017). Vague Credence. *Synthese*, 194 (10), 3931–3954.
- Maher, P. (2006). Review of David Christensen. Putting Logic in its Place: Formal Constraints on Rational Belief. *Notre Dame Journal of Formal Logic*, 47, 133–149.
- Mahtani, A. (2016). Imprecise Probabilities and Unstable Betting Behaviour. *Noûs*.
- Palatini, P. (2009). Recommendations on How to Test Resting Heart Rate. *Medicographia*, 31 (4), 414–420.
- Rinard, S. (Forthcoming). Imprecise Probability and Higher Order Vagueness. *Res Philosophica*.
- Sturgeon, S. (2008). Reason and the Grain of Belief. *Noûs*, 42 (1), 139–165.
- Unger, P. (1979). There Are No Ordinary Things. *Synthese*, 41 (2), 117–154.
- van Fraassen, B. C. (1990). Figures in a Probability Landscape. In M. Dunn, & A. Gupta, *Truth or Consequences* (pp. 345–356). Dordrecht: Kluwer.

van Fraassen, B. C. (2006). Vague Expectation Value Loss. *Philosophical Studies*, 127, 483–491.

Weatherson, B. (2008). Decision Making with Imprecise Probabilities. *Manuscript*.

Wheeler, S. C. (1979). On That Which Is Not. *Synthese*, 41 (2), 155–173.

Williams, R. (2014). Decision-Making Under Indeterminacy. *Philosophers' Imprint*, 14 (4).