IMPRECISE PROBABILITY AND HIGHER ORDER VAGUENESS

Susanna Rinard

Abstract: There is a trade-off between specificity and accuracy in accounts of belief. Descriptions of agents in the tripartite account, which recognizes three doxastic attitudes—belief, disbelief, and suspension—are accurate, but not specific. The orthodox Bayesian account, which requires real-valued credences, is specific, but often inaccurate. I argue that a popular attempt to fix the Bayesian account by using sets of functions is also inaccurate; it suffers from a problem analogous to higher order vagueness. Ultimately, I argue, the only way to avoid these problems is to endorse a principle with the surprising consequence that the trade-off between accuracy and specificity is in-principle unavoidable. However, we can nonetheless improve on both the tripartite and existing Bayesian accounts. I construct a new framework that allows descriptions that are much more specific than those of the tripartite account and yet remain, unlike existing Bayesian accounts, perfectly accurate.

1 Introduction

Much traditional epistemology employs a tripartite account of belief, which recognizes three doxastic attitudes: belief, disbelief and suspension of judgment. However, this approach is too coarse-grained. Its descriptions of agents are typically accurate, but they are not sufficiently specific. For example, one may believe both P and Q, and yet be more confident of one than the other. The tripartite account is blind to these differences in confidence, and yet they are crucially important (e.g., in the explanation of action).

Observations like these are often taken to motivate the orthodox Bayesian account, which recognizes uncountably many different doxastic attitudes: one for each real number between 0 and 1. However, this account suffers from the opposite problem. It is too fine-grained. Descriptions of agents in this account are perfectly specific—reflecting even the subtlest differences

in confidence—but this comes at the cost of accuracy. Sometimes we lack precise, point-valued credences.

The tripartite account, then, has accuracy without specificity; the orthodox Bayesian account has specificity without accuracy. One aim of this paper is to present powerful reasons for the claim that this trade-off between accuracy and specificity is in principle unavoidable: it is simply not possible to represent belief in a way that is both fully accurate and maximally specific.[1]

A second aim of the paper is to show that even if this is right, we can still improve on both the tripartite and orthodox Bayesian accounts. I construct a new *minimal* framework for the representation of belief that allows for much greater specificity than the tripartite account and yet remains, unlike the Bayesian account, perfectly accurate.

In section 2, I discuss a popular attempt to improve on the orthodox Bayesian account by using a set of functions to represent an agent, rather than a single function. It turns out that this set-of-functions account suffers from a slightly different version of the same sort of inaccuracy: specifically, it postulates interval-valued credences with perfectly precise endpoints. This is analogous to the problem of higher-order vagueness. We see that different approaches to vagueness inspire analogous accounts of belief, but each suffers from an analog of the higher-order vagueness problem.

In section 3, I argue that the only way to avoid these problems is by endorsing a principle I call Insurmountable Unclassifiability. I then ask what can be said about the representation of belief if this principle is true. As a preliminary to answering this question, in section 4, I present a novel way of using sets of functions, and intervals, to give partial characterizations of doxastic states. This *minimal interpretation* allows us to give multiple accurate characterizations of the same agent at different levels of specificity.

In section 5, I make use of the minimal interpretation in showing that if Insurmountable Unclassifiability is true, then the trade-off between accuracy and specificity in the representation of belief is in principle unavoidable; it is impossible to characterize an agent's doxastic state in a way that is both fully accurate and maximally specific.

Important open questions remain about Insurmountable Unclassifiability; a complete defense of it is beyond the scope of this paper. My aim is simply to present a powerful reason in favor of it—namely, that it is necessary for avoiding inaccuracy due to over precision in the representation of belief (and analogous higher order vagueness problems)—and to begin exploring some of its surprising and radical consequences.

---

[1] A representation is *specific* insofar as it is committal about whether *P*, for propositions *P* about the agent's doxastic state. It is *accurate* insofar as the agent really is as represented.

2 The Vagueness Analogy

Sometimes we lack the perfectly precise confidence levels required by orthodox Bayesianism. For example, consider LUCKY, the proposition that you will find a four-leaf clover tomorrow. It is simply not the case that there is some unique real number—precise down to the millionth decimal place (and beyond)—that best represents your level of confidence in LUCKY. Any particular choice would involve an element of arbitrariness. Moreover, this doesn't reflect a failure of rationality; given your evidence, even a perfectly rational agent may lack numerically precise credences.

One popular attempt to fix this problem involves using a set of functions to represent each agent, rather than a single function. (Proponents include Jeffrey [1983], van Fraassen [1990], and Joyce [2005; 2010].) This set generates interval-valued credences; one's credence in *P* is the interval that contains all and only the real numbers *r* such that some function in one's set has Pr(P) = *r*.[2]

However, this approach faces a problem much like the one that undermined the single-function approach. The problem of arbitrariness resurfaces, just in a slightly different form. What, exactly, is the upper endpoint of your interval-valued credence for LUCKY? 0.0001? 0.00009? 0.000121? Again, any particular number seems arbitrary.[3, 4, 5] We can gain a deeper understanding of these problems of arbitrariness by pursuing an analogy with vagueness. I begin by defining a predicate MC. Roughly, MC applies to a number just in case it is greater than your level of confidence in LUCKY. For a more precise definition, first, let B[r] say that a coin with bias *r* toward heads will land heads on the next toss. For example, B[0.7] says that a coin with bias 0.7 toward heads will land heads on the next toss. Assume that, in accordance with rationality, for all *r*, your credence in B[r] is *r*. Then we can define MC (for *more confident*) as the predicate that applies to a number *r* just in case you're more confident of B[r] than LUCKY.

---

[2] Some (e.g., Joyce [2010]) allow for non-convex sets of functions, and thus allow for set-valued credences that do not form intervals. For simplicity, I will assume in the text that the set forms an interval, but nothing important hangs on this (see also footnote 4).

[3] The problem of interval endpoint arbitrariness is also raised (though not under that name) in Sturgeon 2008, 158, and Maher 2006, and discussed in Kaplan 2010.

[4] The problem arises even if one's set-valued credence doesn't form an interval (see footnote 1). Any such set must have a unique, numerically precise least upper bound and a greatest lower bound. This is arbitrary in precisely the same way as interval endpoints.

[5] In some cases, it may seem that there are non-arbitrary interval endpoints. For example, suppose your only information is that the chance of drawing a green marble is between 0.2 and 0.4. Then, isn't it rational to have an interval-valued credence whose endpoints are precisely 0.2 and 0.4? In fact, I have argued elsewhere that this is not the rational attitude to take in these cases (see Rinard 2013). However, for our purposes here, it doesn't matter; even if there are some cases in which there are non-arbitrary interval endpoints, the point is that in many, probably most, cases, this is not so. It is the latter cases that are our focus here.

MC is analogous to paradigm vague predicates like TALL. In each case there is a spectrum such that the predicate clearly applies at one end, and clearly fails to apply at the other, but there does not seem to be a sharp boundary. Both seem to admit of borderline cases, and give rise to sorites-style paradoxes.

Importantly, there is a close connection between our lack of a precise credence in LUCKY and the vagueness of MC. The only way to have some precise credence c in LUCKY is for MC not to be vague, that is, for there to be a sharp boundary between MC and not-MC: for MC to be true of all numbers greater than c, and false of all numbers equal to or less than c. A natural thought, then, is that one source of credal imprecision is the vagueness of predicates like MC.

Since the set-of-functions proposal is an attempt to account for credal imprecision, given the close connection between credal imprecision and vagueness it is not surprising that, on one interpretation—found, for example, in van Fraassen 1990, 2005, 2006 and Hájek 2003—the set-of-functions proposal is seen as an instance of the most popular theory of vagueness: supervaluationism. On this interpretation, each function in your set is one admissible precisification of your doxastic state. Functions excluded from your set are inadmissible precisifications. A proposition about your doxastic state is determinately true if true according to all functions in your set. It's indeterminate if true according to some, but not all, functions in your set.

For example, if all functions in your set have Pr(B[0.9]) > Pr(LUCKY), it's determinate that you're more confident of B[0.9] than LUCKY, i.e., it's determinate that MC(0.9) is true. If some functions in your set have Pr(B[0.0001]) > Pr(LUCKY) while others do not, it's indeterminate whether you're more confident of B([0.0001]) than LUCKY, i.e., it's indeterminate whether MC(0.0001) is true. So, on the supervaluationist interpretation, we can see the set-of-functions account as an attempt to account for the vagueness of MC by introducing a third category, indeterminately MC, just as the supervaluationist tries to account for the vagueness of TALL by introducing a third category, indeterminately TALL.[6]

Supervaluationism faces a well-known problem: it requires a sharp boundary where, intuitively, there shouldn't be—namely, between the determinate and indeterminate. This is the so-called problem of higher-order vagueness. On the supervaluationist interpretation of the set-of-functions account of doxastic states, the problematic requirement of precise endpoints for interval-valued credences is just an instance of this more general problem. A sharp upper endpoint for one's interval-valued credence in LUCKY would constitute a sharp boundary between the numbers to which

---

[6] Supervaluationist approaches to vagueness are defended in Fine 1975 and Keefe 2000, among others. Other "third-category" non-supervaluationst approaches, such as Richard 2010, would also generalize to vague doxastic predicates.

MC determinately applies and those to which it indeterminately applies. But there is no sharp boundary here.[7]

As the literature has shown, it is extremely difficult to do justice to higher order vagueness. Below I briefly review a couple of prominent attempts and why they are unsatisfactory. In each case, there is an analogous account of belief that faces analogous problems.

First, some postulate an infinite hierarchy of borderline cases, borderline borderline cases, etc. For example, indeterminately tall heights are first-order borderline cases, but there are also second-order borderline cases: heights that are borderlines of the distinction between determinately tall and indeterminately tall. For every natural number *n*, there are *n*th-order borderline cases. This strategy could be employed in an attempt to fix the problem of precise endpoints for interval-valued credences. We would account for the lack of a sharp line between determinately MC and indeterminately MC by postulating borderline cases of that distinction, and so on up the hierarchy.

However, there is a compelling objection to this approach. Consider paradigm tall heights that aren't borderline cases at *any* level. Call these *absolutely* tall. If every height must be classified as either absolutely tall, or, for some natural number *n*, a borderline case of *n*th order, (or absolutely not tall), then there will be a sharp cut-off between those that are absolutely tall and those that aren't. But there should not be a sharp cut-off here.[8] Similarly, if every real number between 0 and 1 must be classified as either absolutely MC, or absolutely not MC, or, for some natural number *n*, a borderline case of *n*th order, then there will be a sharp cut-off between the numbers that are absolutely MC and those that are not. But there is no sharp boundary here.

A different approach to vagueness involves multiple *degrees* of truth.[9] Some postulate uncountably many: one for each number between 0 and 1, inclusive. This view also faces the problem of higher order vagueness. If we assign some precise degree of truth to every proposition of the form TALL(r), there will be a sharp cut-off between the numbers such that TALL(r) is true to degree 1, and the numbers such that TALL(r) is true to some degree less than 1. But there should not be a sharp cut-off here.

The analogous account of belief, on which each proposition of the form MC(r) is assigned some precise degree of truth, suffers from the same problem; it requires a sharp cut-off between those numbers for which MC(r) is true to degree 1, and those for which MC(r) is true to some degree less than 1. But again, there is no sharp boundary here.

---

[7] Although the supervaluationist interpretation does suffer from this problem, it avoids certain other problems. In Rinard 2015, I defend the supervaluationist set of functions model against the often-made claim that it leads to decision-theoretic incoherence.

[8] Similar problems are presented in Sainsbury 1991 and Wright 2009.

[9] See, for example, Smith 2008 and Zadeh 1975.

So far I have described problems with existing accounts of vagueness, and how those problems survive generalization to the doxastic realm. But the analogy between doxastic predicates like MC, and paradigm vague predicates, goes both ways. As mentioned above, one interpretation of the set-of-functions approach sees it as an instance of supervaluationism. But there is an alternative interpretation, found, for instance, in Schoenfield 2012 and Kaplan 2010. This alternative approach would generalize to paradigm vague predicates; but, in both cases, the approach suffers from a version of the same problem facing other accounts—namely, commitment to a sharp line where there shouldn't be any.

Like the supervaluationist interpretation, on this alternative interpretation, if all functions in your set have $Pr(A) > Pr(B)$, then it's determinate that you're more confident of A than B. However, unlike the supervaluationist interpretation, on this alternative, if some functions in your set have $Pr(A) > Pr(B)$, while others have $Pr(A) < Pr(B)$, and still others $Pr(A) = Pr(B)$, then it's false that you're more confident of A than B, false that you're less confident of A than B, and false that you're equally confident of A than B. Your attitude toward A and B falls into a fourth category. (This is analogous to a view in value theory defended, for example, in Chang 2002, on which it can be the case that A is neither more valuable than B, nor less valuable, nor equally valuable.) A virtue of this view is that it does not require an abrupt transition between cases in which the agent is more confident of A than B, and cases in which the agent is either equally confident, or less confident of A than B. The problem, however, is that it *does* require a sharp line between cases in which the agent is more confident of A than B, and cases in which the agent is *not* more confident of A than B—and, thus, a sharp line between MC and not-MC.

An analogous view for a paradigm vague predicate, such as TALL, would classify so-called borderline cases as neither TALL nor SHORT. Problematically, though, it would require a sharp line between TALL and not-TALL (into which the neither-tall-nor-short cases would fall).

3 Insurmountable Unclassifiability

As we have seen, the orthodox Bayesian account of belief is inaccurate. The nature of the inaccuracy is that it requires point-valued credences, which carries a commitment to the existence of sharp lines where, intuitively, there aren't any—such as between MC and not-MC. (MC(r) says, recall, that you are more confident of B[r]—that a coin with bias r toward heads will land heads on the next toss—than LUCKY [that you will find a four-leaf clover tomorrow].)

So, we cannot give an exhaustive classification of all numbers between 0 and 1 using just the two categories MC and not-MC. In the previous section, we surveyed a number of different attempts to reach a categorization that *is* genuinely exhaustive by refining this initial two-way distinction.

Specifically, each attempt involved adding more (in some cases, infinitely more) intermediate categories. But all these attempts foundered on the same sort of problem: each ended up requiring sharp lines where, intuitively, there shouldn't be any.

In fact, we can identify a more substantive commonality among the sources of failure of these different attempts. In each case, the proposed account failed because it postulated a sharp line between those numbers r such that MC(r) is as true as any proposition possibly could be, according to that account, and those such that this fails to be the case. According to the supervaluationist, for example, for every r, MC(r) falls into one of these three categories: determinately true, indeterminate, and determinately false. This amounts to a sharp line between those numbers such that MC(r) is determinately true—which is as true as any proposition possibly could be, according to this approach—and those numbers for which this fails to be the case. Postulating an infinite hierarchy of borderline cases doesn't help. We would then have a sharp line between the numbers such that MC(r) is *absolutely* true (not borderline at any level)—which is as true as any proposition possibly could be, according to this approach—and those for which this fails to be the case (i.e., when MC(r) is either a borderline case at some level, or absolutely false). Many-valued logics require a sharp line between those numbers for which MC(r) is true to degree 1, and those true to some degree less than 1.

The crucial failing shared by these accounts, then, is their joint commitment to a sharp line between the numbers r such that MC(r) is as true as any proposition possibly could be, and those for which this is not the case. If we want to avoid this problem, then we are forced to deny that there is a sharp line here. (Sainsbury [1991] makes a similar point.) This commitment can be made more precise as follows:

> **Insurmountable Unclassifiability:** It is impossible to correctly classify every number between 0 and 1 into any set of categories C with the following two properties: (1) one of the categories in C is such that, if a number r falls into that category, then MC(r) is as true as possible (2) all other categories in C are such that, if a number r falls into one of them, then it's not the case that MC(r) is as true as possible.[10,11]

Insurmountable Unclassifiability is intended to capture the idea that not only can we not classify all numbers along the two-way MC/not-MC distinction, but we also cannot classify all numbers according to any set of

---

[10] This principle can be generalized to any other vague predicate V by replacing MC with V throughout, and replacing the set of real numbers between 0 and 1 with the set of all instances along the dimension varied in the corresponding sorites argument. For example, for RED it will be the set of all shades; for HEAP, the set of all possible numbers of beans; etc.

[11] This is similar to Terry Horgan's transvaluationist approach. For example, see Horgan 1998, 2010.

categories that is intended to *supersede* this two-way distinction or *refine* it via the addition of (even infinitely many) more intermediate categories. What is wrong with the initial two-way distinction is not that there are too few categories into which to sort numbers. Rather, the fundamental problem—which is shared by refinements of it— is that every category (no matter how many there are) is committal with respect to whether or not the numbers in that category are such that MC(r) is as true as possible. What Insurmountable Unclassifiability denies is that all numbers between 0 and 1 are classifiable into any such set of categories.[12]

Now, it is clear that *some* numbers are classifiable in categories that are members of a set with properties (1) and (2). There are some numbers such that MC(r) is as true as possible (e.g., 0.8, 1, 0.7) and some such that it's not (e.g., 0). What Insurmountable Unclassifiability denies is just that *all* numbers are so classifiable. That some, but not all, numbers are so classifiable might lead one to think that we can draw a sharp line between those that are and those that aren't. But, as we will see, this, too, is not possible. That it's not begins to illustrate the truly *insurmountable* nature of the unclassifiability.

Suppose we were able to say for each number whether it was so classifiable or not. If so, then, for each classifiable number *r* we can say whether or not MC(r) is as true as possible. (If we couldn't, then *r* wouldn't be in the classifiable category.) So, there would be a sharp two-way distinction between, on the one hand, the numbers that are classifiable, and such that MC(r) is as true as possible, and, on the other hand, those that are either unclassifiable, or, classifiable, but not as such that MC(r) is as true as possible. But that would constitute a sharp two-way classification between the numbers for which MC(r) is as true as possible, and those for which it is not! Each number in the first category—classifiable, and such that MC(r) is true as possible—is clearly such that MC(r) is true as possible. Of the numbers in the second category, those that are classifiable, but not such that MC(r) is true as possible, are, obviously, not such that MC(r) is true as possible. And—crucially—those in the second category that are not classifiable are such that it's not the case that MC(r) is true as possible. After all, if MC(r) *were* as true as possible, then that number would fall into the first category. So Insurmountable Unclassifiability entails that, although

---

[12] Those with supervaluationist inclinations might think the solution is to posit another category, consisting of borderline cases of the distinction between what's as true as possible and what's not. But there can be no borderlines of this distinction. Suppose there were. Then there would be a distinction between the determinately true as possible and the borderline true as possible. Consider an instance of the second category. Since it is in the second category, rather than the first, it fails to be as true as it possibly could be. (It could be in the first category.) So what we would have is not *borderline* as true as possible, but rather, determinately *not* as true as possible. It is part of the nature of the distinction between what's as true *as possible*, and what's not, that there can be no borderlines of it. (Compare the argument in Broome 1997 for his Collapsing Principle and an argument in Barnes 1982, 55.)

some numbers are classifiable, and it's not the case that all numbers are classifiable, we cannot categorize each number as either classifiable or not.

Before concluding this section, I'll make an important clarificatory remark. I have spoken frequently of whether it is "possible to classify" a number in such-and-such a way. Such language has the potential to be misleading, insofar as it may suggest that the claims I'm making primarily concern limitations of our abilities. But I do not mean such language to be understood in this way. By "*r* is classifiable as ϕ" I mean nothing more, and nothing less, than "it is the case that *r* is ϕ" Consequently, Insurmountable Unclassifiability is not properly understood as a claim about our abilities, but rather as a substantive metaphysical claim about the nature of the doxastic realm.

Accepting this claim, I have argued, is the only way to avoid commitment to counterintuitive sharp lines. This constitutes a powerful reason in favor of accepting Insurmountable Unclassifiability.

However, there remain important open questions about Insurmountable Unclassifiability. For example, if it is true, then an alternative to classical logic is required (one that, for example, does not vindicate the law of excluded middle). A full defense of the principle would involve constructing a new account of vagueness that addresses these issues. Doing so is beyond the scope of this paper; however, my aim is simply to establish that, if we want to avoid problems of higher-order vagueness, then we must consider Insurmountable Unclassifiability a desideratum for any adequate account of vagueness.

Having presented some powerful reasons in its favor, I turn now to exploring some of the surprising and radical consequences Insurmountable Unclassifiability has for the representation of doxastic states.

4 The Minimal Interpretation

I began this paper by noting a trade-off between accuracy and specificity in two popular accounts of belief. In the section following this one, I will show that if Insurmountable Unclassifiability is true, then this trade-off is in principle unavoidable; it is impossible to characterize belief in a way that is both fully accurate and maximally specific. In order to do this, I will first present a new way of using interval notation and sets of functions to represent belief, which I call the *minimal interpretation*. That is the aim of this section.

First, notice that among informal characterizations of belief, some are more specific than others. For example, suppose I say you're *more confident of P than not*, and then elaborate that you're *nearly certain of P*. The second description is more specific than the first, but both are perfectly accurate. The minimal interpretation of sets and intervals presented here allows us to give, in a similar fashion, multiple accurate characterizations, at different levels of specificity, of a single agent's doxastic state.

I'll start with intervals. First I'll give a definition that is helpful, but potentially misleading; then I'll define it more carefully. Helpful, but misleading: on the minimal interpretation, [c, d] accurately characterizes an agent's doxastic attitude toward *H* just in case *the agent's level of confidence in H* is contained within [c, d]. Note that on this interpretation, there will always be multiple accurate intervals at varying levels of specificity. On reason is that, if [c, d] is accurate, then any larger (more inclusive) interval is also accurate. This is because, if one's level of confidence is contained within [c, d], then it is contained within any interval of which [c, d] is a subset. For example, since I am quite confident that LUCKY is false, my doxastic attitude toward LUCKY is accurately characterized by [0, 0.7], [0, 1], [0, 0.5], and many others. On the minimal interpretation, each of these is perfectly accurate; some are more specific than others.[13]

The definition given above is misleading insofar as it suggests that the agent has a point-valued level of confidence. The fix is: on the minimal interpretation, an interval [c, d] accurately characterizes an agent's doxastic state toward *H* just in case (1) the agent is more confident of *H* than any proposition in which her credence is less than *c*; and (2) the agent is less confident of *H* than any proposition in which her credence is greater than *d*. One final stipulation: the conditions on the right-hand side must be *as true as possible* for the interval to count as accurate.

It is important to note that the minimal interpretation does not, by itself, constitute a substantive claim about the nature of doxastic states. It is just a new way of using notation. One way to see this is by noting that the minimal interpretation is entirely compatible with the substantive views held by proponents of different views on credal imprecision, including supervaluationism, an infinite hierarchy of borderline cases, many-valued logics, Insurmountable Unclassifiability, even the single-function account, etc. For example, if the supervaluationist has [c, d] as the agent's unique interval-valued credence in *H*, then that interval (as well as any more inclusive interval) automatically counts as accurate on the minimal interpretation. This is because the supervaluationist interpretation of the interval notation is logically stronger than the minimal interpretation of the same notation (hence the name *minimal*). For another example, if one has a point-valued credence *c* in *H*, then every interval of which *c* is a member will count as accurate on the minimal interpretation. In general, *to characterize an agent's doxastic attitude using the new, minimal interval notation is to remain neutral on the issues that divide defenders of all different views on credal (im)precision.* This is because, on the minimal interpretation, to characterize an agent's attitude toward *H* with some interval is to remain completely neutral about the status of numbers inside that interval. For any such number *r*, it may be that one is more confident of *H* than B[r]; or less

---

[13] Walley (1991) discusses related proposals under the category of "non-exhaustive" interpretations.

confident, or equally confident, and it may be indeterminate whether *r* is one's level of confidence in *H*. For any possible status *r* might have, we remain completely neutral about whether *r* has that status. To describe one's attitude with an interval is to be committal only about numbers outside that interval.

We can give a new, minimal re-interpretation of sets of functions in the same vein. On a first pass, a set counts as accurate just in case every proposition about the agent's doxastic state that is true according to all functions in that set is as true as possible of the agent. This is only a first-pass definition, though, because some propositions true according to all functions in the set are not true of the agent. For example, it is true according to every function in the set that you have a precise credence in LUCKY. (The functions agree that you have a precise credence; they just disagree about what it is.) But *that you have a precise credence in LUCKY* is precisely what we want to deny! We can get to the root of this problem by noticing that, although the existential claim *there is some real number r such that r is your precise credence in LUCKY* is true according to every function, there is no instance of that existential claim that is true according to every function—that is, there is no *r* such that *r is your precise credence in LUCKY* is true according to every function in the set. So, we can fix the problem by revising the definition as follows. First, let *Z* be the set of all propositions true according to all functions in the set. Generate *Z*- by removing from *Z* any proposition that is, or is equivalent to, some existential claim such that no instance of that existential claim is true according to every function in the set. Now the proper definition: on the minimal interpretation, a set counts as accurate just in case every proposition in *Z*- is true (as possible) of the agent.

To characterize an agent with a set, on this new, minimal interpretation, is to remain entirely neutral on the status of propositions about which different functions in the set disagree, just as the minimal interval notation remains neutral on the status of numbers inside the interval. There will always be multiple accurate sets, just as there are always multiple accurate intervals. In addition, as before, this interpretation is compatible with different views on credal imprecision. For example, if *S* is the agent's unique set of functions, interpreted in the supervaluationist way, then *S* (and any more inclusive set) counts as an accurate description of that agent on the minimal interpretation. If an agent is best represented by a single credence function, any set containing that function is accurate.

This minimal interpretation can accommodate traditional Bayesian approaches to a wide range of different issues, such as learning from experience, theory confirmation in science, and decision theory. It will typically be the case that if, on the single-function account, some proposition *B* is true of an agent if that agent's credence function has property *Q*, then, on the minimal account, *B* will be true of the agent if there is some set of

functions, which accurately characterizes that agent, all of whose members have property Q.[14] For example, on the minimal account we can say that E confirms H, for agent A, if there is some accurate set of functions, all of whose members have $Pr(H|E) > Pr(H)$. The agent satisfies minimal synchronic requirements for rationality if there is some accurate set of functions, all of whose members satisfy minimal synchronic probabilistic coherence requirements.[15] The agent rationally updates on evidence E, received between $t1$ and $t2$, just in case any set of functions $S2$ that is accurate at $t2$ can be obtained from some set $S1$ that was accurate at $t1$ via conditionalizing each function in $S1$ on E, and so forth.

What I want to emphasize here, though, is how, with respect to the representation of belief, the minimal interpretation constitutes a tool that allows us to improve on both the tripartite and existing Bayesian approaches. We can use it to give characterizations that are much more specific than those of the tripartite account. For example, if [0.95, 0.96] accurately characterizes my attitude toward A, and [0.97, 0.98] my attitude toward B, this represents (among other things) that I am more confident of B than A, even if I believe both. Yet, unlike those of existing Bayesian accounts, these characterizations remain perfectly accurate. They do not fall prey to the problem of arbitrariness, since accurate intervals are not taken to be *uniquely* accurate. It is perfectly accurate to characterize my attitude toward LUCKY with [0, 0.8], even though the endpoints are perfectly precise, because this characterization is not taken to be a unique best one. Other intervals, such as [0, 0.7] or [0, 0.6], may be equally accurate.

The only shortcoming of the minimal interpretation is that it facilitates only partial descriptions of the agent's doxastic state. To characterize one's doxastic state with [0, 0.7], for example, is not maximally informative, since narrower intervals may be equally accurate. Whether this shortcoming can be remedied is considered in the next section.

5 The Combination of Perfect Accuracy and Maximal Specificity is Impossible

In the previous section, I presented the minimal interpretation, which allows us to give multiple accurate descriptions of an agent's doxastic state at different levels of specificity. It is fine to characterize one's doxastic state by giving a few descriptions of this kind, but doing so raises a natural question. What is the *most specific* accurate interval? For example, above I listed the following as accurate characterizations of my attitude toward

---

[14] For reasons similar to those given above, this will not be true when B is an existential claim, none of whose instances is agreed upon by all functions.

[15] If Insurmountable Unclassifiability is true, then, as noted in the previous section, the correct logic is non-classical. This will require us to re-think what probabilistic coherence looks like, given that probability theory is typically developed (axioms stated, theorems proved from those axioms) in a classical environment.

LUCKY: [0, 1], [0, 0.8], and [0, 0.5]. Having done this, it is natural to ask about other intervals. What about [0, 0.3]? (0, 0.27)? [0, 0.2]? In particular, it is natural to wonder which interval is the *most specific* interval that is still accurate. After all, any information encoded in a less specific interval is also contained in a more specific interval—so why not just isolate the maximally specific accurate interval, identify it as such, and forget about the rest?

Interestingly, it turns out that if we are serious about avoiding counterintuitive sharp lines—i.e, if we accept Insurmountable Unclassifiability—then we must regard this very natural thought as deeply flawed. Suppose there were a most specific interval that accurately characterized my attitude toward LUCKY—say, [c, d]. Then any narrower interval would fail to be accurate, but then d would constitute a precise boundary where there shouldn't be one. For all numbers greater than d, MC(r) would be as true as possible; for all numbers equal to or less than d, MC(r) would fail to be as true as possible. But, as we have seen, Insurmountable Unclassifiability entails that we can't classify all numbers r in this way. The upshot is that if we accept Insurmountable Unclassifiability—a desideratum for any account of vagueness that hopes to avoid problems of arbitrariness and higher-order vagueness—then sometimes *there is no maximally specific, fully accurate characterization of one's doxastic state.*

This is a surprising and radical conclusion. It means that it is impossible to give a complete description of an agent's doxastic state; there is no such thing as a complete description. We can give partial descriptions of an agent's doxastic state; some will be more specific than others. But we can never identify a particular description as *maximally specific*, or *complete*.[16]

6 The Forced March

Consider the following finite series of intervals. All intervals in the series share the same lower endpoint: 0. The upper endpoint of the first interval is 1. For each interval in the series, its upper endpoint is the result of

---

[16] How might this affect other issues, such as updating or confirmation? In a sense, not at all—these are still handled in the minimal account exactly as described in the previous section. But Insurmountable Unclassifiability introduces a new wrinkle: there is now no *guarantee* that there will be a fact of the matter about whether the relevant conditions obtain (although there *may* always be a fact of the matter). For example, on the minimal account we have that *if* there is some accurate set of functions, all of whose members have Pr(H|E) > Pr(H), then *E* confirms *H*. And if there is no accurate set of functions, all of whose members have Pr(H|E) > Pr(H), then it's not the case that *E* confirms *H*. But since, according to Insurmountable Unclassifiability, it's not the case that every set is classifiable as either accurate or not, the door is now open to the possibility that it will not be settled whether *E* confirms *H*. It is important to emphasize, however, that Insurmountable Unclassifiability does not *require* that there sometimes fail to be a fact of the matter here. It is compatible with Insurmountable Unclassifiability that there is always a fact of the matter about whether *E* confirms *H*. The same goes for other issues, such as rational updating.

subtracting some minuscule positive real number ϵ from the upper endpoint of its predecessor. The series ends when we reach a number that is equal to or less than 0.

Imagine going through the intervals of this series one by one and asking someone whether it is accurate concerning their level of confidence in LUCKY. All sorts of things might happen, of course: the speaker might get annoyed and walk off; they might try to change the subject; etc. Philosophers are not usually interested in these scenarios, however. Rather, it is typically assumed that the speaker in question has various features that rule out such possibilities. It is important to begin any discussion of the forced march by making such assumptions perfectly clear and explicit.

Here are some stipulations that are commonly made:

(1) For each question, if the answer is in fact "yes," then the speaker knows that it is.
(2) If the speaker knows that the answer to a question is "yes," then, when that question is asked, they will say "yes."
(3) The speaker will say "yes" *only* if they know that the answer is "yes."
(4) For every question, there is a fact of the matter about whether or not the speaker answers "yes."[17]

Any one of these stipulations, on its own, is fairly innocent. But to stipulate that they hold all at once is a highly theoretically loaded act, for that amounts to stipulating that Insurmountable Unclassifiability is false. From (1) and the factivity of knowledge we have that, for each question, the answer is "yes" iff the speaker knows that it is; from (2) and (3), that they know the answer is "yes" iff they say "yes." It follows that the answer is "yes" iff the speaker says "yes." Given this, stipulation (4)—that there's always a fact of the matter about whether the speaker says "yes"—is equivalent to the stipulation that, for every question, there's a fact of the matter about whether "yes" would be a correct answer, i.e., that for every interval in the series, there's a fact of the matter about whether or not it's accurate. But the denial of this is a core commitment of Insurmountable Unclassifiability.

If Insurmountable Unclassifiability is right, then, there is no possible situation in which (1)–(4) hold all at once. However, for any proper subset of {(1), . . . (4)}, there is a possible situation in which the assumptions in that subset hold. For example, suppose we retain (1)–(3), but give up (4). Then, the speaker may answer the series of questions by emitting a series of sounds which gradually morph from "yes" to "no." The first few are clearly "yes," and the last are clearly "no," but adjacent sounds in the series are indistinguishable. If the speaker's voice is not sufficiently sophisticated to pull off such an act, she may need a prop, such as a collection of colored cards that constitute a sorites series from red to orange. She can stipulate

---

[17] For one representative discussion of the forced march, see Williamson 1994.

that her answer is "yes" iff she holds up a red card, and then proceed to answer the questions by holding up each of the cards in order.

Alternatively, suppose (4) does hold, along with (1) and (3), but not (2). There are many possibilities consistent with this, and with Insurmountable Unclassifiability. For example, the speaker may respond to all questions in a maximally non-committal way: failing to say "yes," but not thereby asserting (or implicating or suggesting) that the answer is not "yes." (Which speech act can accomplish this will depend on the context; silence is a natural candidate, but only if remaining silent is not taken to implicate that "yes" would be incorrect.) Or, the speaker may answer "yes" to the first one or two questions, then give the non-committal response, answer "yes" a few more times, and go back to the non-committal response. Or, she may answer "yes" for a while, and then switch to the non-committal response for a while, and then answer "no" to the last few questions. The point is that the non-committal speech act gives the speaker the option of answering "yes" to one but not the other of two adjacent questions without thereby committing herself to a sharp line of the sort that Insurmountable Unclassifiability denies.

The upshot is that different stipulations make for radically different versions of the forced march scenario. We must take care to make explicit what is, and what is not, being assumed about the particular case under discussion. It is uncontroversial that, for each of (1)–(4) individually, there is some possible situation in which it holds. But it would be a mistake to conclude that it is similarly innocuous to suppose that there is a possible situation in which they all hold together. This claim, far from being an innocent stipulation, is highly theoretically loaded. Indeed, it is tantamount to the denial of Insurmountable Unclassifiability.

7 Conclusion

I began by noting a trade-off between accuracy and specificity in existing accounts of belief. The tripartite account of traditional epistemology has accuracy without specificity. The Bayesian accounts of formal epistemology (both the single-function and set-of-functions version) have specificity without accuracy.

There are powerful reasons for thinking that this trade-off is in-principle unavoidable. The only way to avoid commitment to counterintuitive sharp lines—and the analogous problem of higher order vagueness—is to endorse Insurmountable Unclassifiability. But if this principle is true, then it is simply not possible to characterize belief in a way that is both fully accurate and maximally specific.

What we can do, however, is improve on both the tripartite account and the existing Bayesian accounts of belief. The new framework constructed here enables us to use intervals, and sets of functions, to give partial descriptions that are much more specific than the tripartite account, and

yet which, unlike the Bayesian accounts, do not fall prey to problems of over-precision and so are perfectly accurate.

Susanna Rinard
E-mail : susannarinard@fas.harvard.edu

References:

Barnes, Jonathan. 1982. "Medicine, Experience, and Logic." In *Science and Speculation*, edited by Jonathan Barnes, Jacques Brunschwig, Myles Burnyeat, and Malcolm Schofield, 24–65. Cambridge: Cambridge University Press.

Broome, John. 1997. "Is Incommensurability Vagueness?" In *Incommensurability, Incomparability, and Practical Reason*, edited by Ruth Chang, 67–89. Cambridge: Harvard University Press.

Chang, Ruth. 2002. "The Possibility of Parity." *Ethics* 112 (4): 659–688. http://dx.doi.org/10.1086/339673.

Fine, Kit. 1975. "Vagueness, Truth and Logic." *Synthese* 30 (3–4): 265–300. http://dx.doi.org/10.1007/BF00485047.

Hájek, Alan. 2003. "What Conditional Probability Could Not Be." *Synthese* 137 (3): 273–323. http://dx.doi.org/10.1023/B:SYNT.0000004904.91112.16.

Horgan, Terence. 1998. "The Transvaluationist Conception of Vagueness." *Monist* 81 (2): 313–330. http://dx.doi.org/10.5840/monist19988122.

Horgan, Terence. 2010. "Transvaluationism about Vagueness: A Progress Report." Reprinted in Horgan (2017b). *Southern Journal of Philosophy* 48 (1): 67–94. http://dx.doi.org/10.1111/j.2041-6962.2010.01006.x.

Jeffrey, Richard. 1983. "Bayesianism with a Human Face." In *Minnesota Studies in the Philosophy of Science*, edited by John Earman, Vol. 10: Testing Scientific Theories, 133–156. Minneapolis: University of Minnesota Press.

Joyce, James. 2005. "How Probabilities Reflect Evidence." *Philosophical Perspectives* 19 (1): 153–178. http://dx.doi.org/10.1111/j.1520-8583.2005.00058.x.

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

Kaplan, Mark. 2010. "In Defense of Modest Probabilism." *Synthese* 176 (1): 41–55. http://dx.doi.org/10.1007/s11229-009-9483-2.

Keefe, Rosanna. 2000. *Theories of Vagueness*. Cambridge: Cambridge University Press.

Maher, Patrick. 2006. "Book Review: David Christensen: Putting Logic in its Place: Formal Constraints on Rational Belief." *Notre Dame Journal of Formal Logic* 47 (1): 133–149. http://dx.doi.org/10.1305/ndjfl/1143468316.

Richard, Mark. 2010. "Indeterminacy and Truth Value Gaps." In *Cuts and Clouds: Vagueness, its Nature and its Logic*, edited by Richard Dietz and Sebastiano Moruzzi. Oxford: Oxford University Press.

Rinard, Susanna. 2013. "Against Radical Credal Imprecision." *Thought* 2 (1): 157–165.

Rinard, Susanna. 2015. "A Decision Theory for Imprecise Probabilities." *Philosophers' Imprint* 15 (7): 1–16.

Sainsbury, Mark. 1991. "Is There Higher-Order Vagueness?" *Philosophical Quarterly* 41 (163): 167–182. http://dx.doi.org/10.2307/2219591.

Schoenfield, Miriam. 2012. "Chilling Out on Epistemic Rationality." *Philosophical Studies* 158 (2): 197–219. http://dx.doi.org/10.1007/s11098-012-9886-7.

Smith, Nicholas. 2008. *Vagueness and Degrees of Truth*. Oxford: Oxford University Press.

Sturgeon, Scott. 2008. "Reason and the Grain of Belief." *Noûs* 42 (1): 139–165.

van Fraassen, Bas. 1990. "Figures in a Probability Landscape." In *Truth or Consequences*, edited by Michael Dunn and Anil Gupta, 345–356. Dordrecht: Kluwer.

van Fraassen, Bas. 2005. "Conditionalizing on Violated Bell's Inequalities." *Analysis* 65 (1): 27–32. http://dx.doi.org/10.1093/analys/65.1.27.

van Fraassen, Bas. 2006. "Vague Expectation Value Loss." *Philosophical Studies* 127 (3): 483–491. http://dx.doi.org/10.1007/s11098-004-7821-2.

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

Williamson, Timothy. 1994. *Vagueness*. London: Routledge.

Wright, Crispin. 2009. "The Illusion of Higher-Order Vagueness." In *Cuts and Clouds: Vagueness, its Nature and its Logic*, edited by Richard Dietz and Sebastiano Moruzzi, 523–549. Oxford: Oxford University Press.

Zadeh, Lotfi. 1975. "Fuzzy Logic and Approximate Reasoning." *Synthese* 30 (3–4): 407–428. http://dx.doi.org/10.1007/BF00485052.