How to Read a Representor

Edward J. R. Elliott

School of Philosophy, Religion and History of Science
University of Leeds

Abstract

Imprecise probabilities are often modelled using sets of probability functions, known as *representors*. In the recent literature, two ways of interpreting representors have emerged as especially prominent: vagueness interpretations, according to which each probability function in the set represents how the agents beliefs would be if any vagueness were precisified away; and comparativist interpretations, according to which the set represents those comparative confidence relations that are common to all functions therein. I argue that both of these interpretations come with significant limitations. I also propose an alternative interpretation, the functional interpretation, according to which representors are best interpreted by reference to the roles they play in the theories that make use of them.

Keywords: Imprecise probabilities · Representors · Supervaluationism · Comparative probability · Degrees of belief · Functionalism

Introduction

As a model for beliefs, probability functions are great. They might have some issues here and there, to be sure, but there's just so much they get *right*. So I'm a big fan. But you know what's even better than a probability function? A whole bunch of probability functions! Anything that can be represented with probability functions can be represented with sets thereof, plus more besides. So if we switch from modelling our beliefs with probability functions over modelling them with sets of probability functions—what many in philosophy call *representors*, and others call *credal sets*—then it looks like we've got nothing to lose.

Well, maybe something. Aside from some surface-level agreement that representors represent 'imprecise probabilities', and frequent appeals to a *credal committee* metaphor that (as we'll see soon enough) is as apt to mislead as it is to illuminate, there really isn't a great deal of consensus on what it is exactly that representors are supposed to represent nor how they represent it. Worse, while everyone seems to agree that not all of the information built into a representor model need reflect something *real*—read: a genuine property of our belief systems, as opposed to a mere artefact of the formalisation—competing interpretations of the model differ significantly regarding which aspects of a representor should be taken seriously and which are mere artefacts, and hence they differ in non-trivial ways regarding what kinds of inferences can be rightly drawn about an agent's beliefs from a representor model of those beliefs. A better recipe for confusion you will not often find.

There are four sections to this paper. After laying out some more background on representors in §1, I will outline and critically discuss two ways of reading representors that have become prominent in the recent literature: *vagueness interpretations* in §2, and *comparativist interpretations* in §3. Both have important limitations. Finally, in §4, I will present and defend an alternative interpretation—what I'll call the *functional interpretation*.

A note before going further. There are many interpretations I don't discuss here—for instance, that each function in a representor represents a way your beliefs could be consistent with your limited introspective evidence, and hence the set represents your higher-order uncertainty about your own uncertainty. I've chosen to focus my discussion on the vagueness and comparativist interpretations for two reasons. First, an exhaustive taxonomy and evaluation of every conceivable interpretive possibility would make for very tedious reading indeed. Second, and moreover, the point of the discussion is *not* to say that any interpretation other than the one I're proposing is erroneous; rather, it is to set up an illuminating contrast between my interpretation versus some common alternatives. Better to understand what my proposal is and why it's worth considering when you can easily compare it with what it's not.

It would be a gross error, therefore, to read me as arguing that there's only one correct way to use a representor to represent something-to-do-with-our-doxastic-states. No one's making that mistake—of course there's more than one possible legitimate and useful interpretation of a set of probability functions! And it would be just as severe an error to suppose that my argumentative strategy requires conclusively ruling out any semblance of the vagueness and comparativist interpretations. My goal is to prop up the interpretation I like, not demolish the interpretations I don't. It is enough that I highlight concerns and limitations for the more common interpretations which do not arise for the functional interpretation, and hence provide reasons for giving the functional interpretation a bit of time in the limelight as well.[1]

1. Background

We can think of representors as a response to concerns arising for the traditional single-function model of belief. Say that a proposition is a set of possible worlds, and let $\mathbf{P} = \{p, q, r, \ldots\}$ contain all and only those propositions regarding which our subject—Sally—has beliefs to some degree or other. We assume that $\mathbf{P}$ is closed under (relative) complements and intersections, and thus also under unions. Then, according to the traditional model, Sally's degrees of belief can be represented using a single measure, $\mu : \mathbf{P} \mapsto [0, 1]$, which satisfies the usual probabilistic constraints—i.e., for any p, q in $\mathbf{P}$, if p is true at all worlds then $\mu(p) = 1$, and $\mu(p \cup q) = \mu(p) + \mu(q)$ whenever $p \cap q$ is false at all worlds.

[1] It is an amazing thing that, even among those whose work frequently deals with matters of partial evidence and incremental confirmation, still a philosophical argument is considered problematic if it does not provide overwhelming and incontrovertible reasons for accepting the position it's intended to support to the exclusion of any others—one either says enough to convince the sceptic (in the space of a few pages no less!), or they'll have failed to provide a 'good' argument. I'm not going to convince the sceptic. That's a fool's errand.

There is something strikingly unrealistic about the traditional model. We need not go into all of the concerns that have been raised, for they are many and various—see (Jeffrey 1983), (Seidenfeld 1988), (van Fraassen 1990), (Walley 1991), (Kaplan 1996), (Joyce 2005; 2010), (Sturgeon 2008), (Hájek 2012), (Alon and Lehrer 2014), and (Bradley 2014). But it won't hurt to consider one example (adapted from Fishburn 1986). Imagine that before you sits an old pack of cards. You've been told that some of the cards are missing, but that's all you're told—you don't know how many are missing nor which ones. Now consider:

$p =$ The global population in 2100 will be greater than 12 billion
$q =$ The next card drawn from this old deck will be a heart

If you're like most people then you'll have at least some positive degree of confidence towards both p and q, and you almost certainly won't have *exactly* as much confidence in p as you do in q. Given that, here's a question: *how much more or less confident are you in* p *than you are in* q*?*

The traditional model implies that there will be a unique real value r such that you're exactly r times as confident in p as you are in q. And that's clearly implausible. Not because there may be some facts about the strengths of your beliefs that are introspectively hard to determine, though that may be true as well. The problem is simply that such precise values aren't realistically applied to measure a squishy psychological quantity like strength of belief, at least for people like us. Whatever it is about us that grounds the facts about our degrees of belief, there just isn't sufficient information down there to determine that we believe p down to the n^{th} degree for absurdly large n. Indeed, it's not even clear that you must have *more* or *less* confidence in one of p or q over the other. Maybe there's no determinate fact of the matter as to which one you believe more; or maybe they're determinately incomparable. You might even be sympathetic to the idea that your strength of belief in p can be *on a par* with your strength of belief in q, where the *parity* relation is a special kind of symmetric comparative strength relation that holds only if p and q are not believed to exactly the same degree (cf. Chang 2002). But one cannot represent any of these ideas on the traditional model, according to which if p and q belong to $\mathbf{P}$, then you're either more or less confident in p than you are in q, to an infinitely precise degree, or you're exactly equally confident in both.

The upshot is that there appears to be something about the way our beliefs are, or a way at least that they might be, that the traditional single-function model just isn't able to capture. Worry not what that something is just yet, worry only that whatever it is it's missing. Maybe it's several things. In any case, a more general model is needed.

Representors to the rescue! On this new and improved model, we should represent systems of belief by means of a (finite or infinite, but always non-empty) set of probability functions, $\mathbf{R} = \{\mu, \mu', \mu'', \ldots\}$, all defined on the same space of propositions $\mathbf{P}$. When the set contains just a single probability function, then it just represents the very same system of beliefs as would have been represented by that function according to the traditional model. On the other hand, when a representor contains multiple functions, then it represents *something else*. What that 'something else' should be is the topic of this paper, and varies from person to person. But the *credal committee* metaphor (originating with Joyce 2010, p. 288) is often used to give the rough idea. Imagine that every μ in $\mathbf{R}$ gets a single vote on what Sally's beliefs are going to be like, and the vote passes if the committee is unanimous. If every μ votes that Sally's confidence in p is greater than her confidence in q, then Sally's confidence in p really is greater than her confidence in q—even if there's no precise value r such that all members of the committee agree that Sally's confidence in p is exactly r times her confidence in q. Similarly, if some members of the committee vote that Sally is more confident in p than she is in q, while others vote that she's more confident in q than she is in p, then $\mathbf{R}$ as a whole represents neither since the committee failed to come to an agreement on the matter.

(Be warned: while this metaphor can be useful for roughly summarising how a representor represents under some interpretations, it can also be misleading. There are many inferences that are implicated by the metaphor that end up being valid under some interpretations but invalid under others. For example, unreflective application of the credal committee metaphor will suggest that if Sally has strictly more confidence in p than she does in q, then every μ in her representor ought to assign a higher value to p than to q. As we'll see, this holds for *some* of the interpretations I consider below, but not all.)

A little bit more notation and terminology will be useful for what follows. For each representor $\mathbf{R}$ we can define its *summary function* like so:

$$\mathcal{R}(p) = \{\mu(p) : \mu \in \mathbf{R}\}$$

That is, for each p in $\mathbf{P}$, $\mathcal{R}(p)$ picks out the set of values that the μ in $\mathbf{R}$ assign to p. In some cases, $\mathcal{R}(p)$ might be an (open or closed) interval; in others, $\mathcal{R}(p)$ may be 'gappy'. I'll mostly be talking about cases where $\mathcal{R}(p)$ is an interval, but this is mostly for convenience. The more important thing to note is that while summary functions can be good for, e.g., describing the spread of values assigned to a proposition by the 'credal committee', a summary function is not just another way of representing a representor. In some cases distinct representors will determine the very same summary function, and so there is a potential loss of information if we simply replace the former with the latter. One is a set of real-valued functions, the other is a set-of-reals-valued function; they can carry different information, and they shouldn't be confused.

2. The Vagueness Interpretation

Suppose one of us points to a cat and say 'look at that cat'. There is presumably some vagueness as to what we're picking out. In the vicinity of the space where we're pointing there will be a multitude of precise cat-like things, cat_1, cat_2, cat_3, …, differing from one another by molecule here or fraction of a whisker there. We're not really referring to any one of them in particular, though neither are we referring to none of them. Rather, it looks like each serves as a potential referent for 'that cat', and it's simply undecided which it should be. So say that each of cat_1, cat_2, cat_3, …, counts as a *precisification* of what we might be referring to when we say 'that cat'—it's the kind of thing we would be referring to if we were to make our language perfectly precise in one or another way. Say also, at least to begin, that anything that's true of all such precisifications is determinately true, whereas if something is true relative to some precisifications and false on others, then it's indeterminate. Call this the *supervaluationist rule*.

According to vagueness interpretations, representors represent vagueness in our degrees of belief, and they do so via the same supervaluationist rule, or something much like it. On the simplest versions, each μ in $\mathbf{R}$ is to be interpreted as representing a precise system of beliefs as per the traditional model, and $\mathbf{R}$ is then taken to contain all and only the precisifications of Sally's beliefs. Given that, roughly, if $\mathcal{R}(p) = [0.34, 0.37]$, then Sally's degree of belief regarding p is said to be "vague over the $[0.34, 0.37]$ interval"—it's determinately between 0.34 and 0.37 inclusive, but for any precise degree within that interval it's indeterminate whether *that* is the degree to which Sally believes p.[2]

There are other ways I could spell out the details. Traditional supervaluationism says that truth is truth-under-all-precisifications, and something that's true on some precisifications but false on others will simply lack a truth-value. Degree-theoretic supervaluationism says that if something's true on all precisifications then it's 100% true, 0% true if it's false on all precisifications, and some middling degree between 0% and 100% otherwise. There can also be variation regarding whether the vagueness is understood to result from some kind of semantic indecision, or is instead a feature of the belief system itself independent of how we talk about it. So there isn't really one vagueness interpretation, but a family of them. These differences shouldn't matter for my purposes. One can find instances of the vagueness interpretation in (van Fraasen 1990; 2006), (Hájek 2003), (Rinard 2015), and (Levinstein 2019). Hájek and Smithson (2012, §3) and Joyce (2010) also present what *could* be read as instances of the vagueness interpretation, at least under some precisifications.[3]

Limitations with the vagueness interpretation

You might worry that the simple vagueness interpretation is a bit too simple—there are some properties that are automatically represented by any probability function whatsoever according to the traditional model, and which are therefore represented by any non-empty set of probability functions, but which we

[2] That isn't *quite* right, though it'll be good enough for my purposes. More accurate would be to say that if $\mathbf{R} = \{\mu_1, \mu_2, \ldots\}$, then it's indeterminate whether μ_1 represents Sally's total system of beliefs, or μ_2 does, and so on.

[3] We could also define another, strictly broader category of interpretation—the *supervaluational interpretations*—characterised by their shared use of the supervaluationist rule. For example, Williams (2014) interprets sets of probability functions as representing the range of attitudes one may rationally take towards an indeterminate proposition. Along similar lines, one might use a representor to represent the rationally permissible precise belief states relative to an agent's evidence, where the facts about rational permissibility are indeterminate. But these interpretations are only superficially similar to what I'm calling the vagueness interpretation. The difference is that the vagueness interpretation is about representing vagueness or indeterminacy *in our degrees of belief*, as opposed to representing vagueness or indeterminacy *in which degrees of belief are rational*.

shouldn't therefore suppose are held by any vague system of beliefs. Every member of the "credal committee" seems to say that Sally has *some* precise degree of belief for every proposition in **P**, for example, but we probably shouldn't want to say that **R** therefore represents Sally as determinately having some precise degree of belief for every proposition even if it doesn't represent her as having any particular precise degree of belief for any proposition. Likewise, every μ in **R** represents that there is a unique real value r such that Sally is r times more confident in p than she is in q, provided she has some positive degree of confidence in both, but this is precisely the sort of thing we were trying to avoid having to say!

It's natural to say that these sorts of things are an artefact of the model, arising from the use of a set of precise functions to represent an imprecise state, and not to be taken seriously. To suppose otherwise smells a bit like what Lewis once called *fanatical supervaluationism*,

> … which automatically applies the supervaluationist rule to any statement whatever, never mind that the statement makes no sense that way. (1999, p. 173)

So let's grant that there might be some restrictions on how we apply the supervaluationist rule when reading a representor (cf. Zynda 2000, p. 49; Rinard 2017, p. 267). For instance, we might reasonably say that **R** represents as determinately true anything that's true according to every μ in **R**, with the exception of those existential claims where no particular instances of that claim are true according to every such μ.

If we're careful about highlighting these kinds of exceptions, then we can avoid the absurd implications of fanatical supervaluationism. Given that, I don't think the foregoing presents a serious problem, simply because the fix is straightforward. But there is another concern in the vicinity which isn't so easily addressed. This concern arises specifically in those cases where the representor contains functions that (according to the traditional model) represent belief states that are vastly different from one another.

Consider again the precisifications of 'that cat'. Each of cat_1, cat_2, cat_3, and so on, all have very precise boundaries, even though *that cat* does not have precise boundaries. So there's at least one respect in which what's true for every precisification of 'that cat' is not true of that cat. Not a problem: we simply don't apply the supervaluationist rule to every statement whatsoever; easy fix. However—and this is the important part—in all the ways that really matter, every precisification of 'that cat' is still a cat. Each precisification walks like a cat, each meows like a cat. If we were to bundle up all the properties that we ordinarily associate with cathood, then each of cat_1, cat_2, cat_3, …, would satisfy at least the very large majority of them—the main exception of course being their perfectly precise boundaries. They all make sense *qua* precisification of 'that cat'. Because of this, not one of cat_1, cat_2, cat_3, …, will be much like a cassowary. Any precise object that's like enough to a cat that it might serve as a precisification of 'that cat' will be so different from a cassowary that it could not serve as a precisification of 'that cassowary'.

So now compare two representors, $\mathbf{R}_{\text{narrow}}$ and $\mathbf{R}_{\text{wide}}$. $\mathbf{R}_{\text{narrow}}$ determines only a narrow spread of values for some proposition p. Let's say: $\mathcal{R}(p) = [0.34, 0.37]$. We have good idea of what Sally would probably be like if she believed p to this or that precise degree within the $[0.34, 0.37]$ interval. It's not perfect, but decision theory gives us a working sense of what the functional roles of such belief states are like, how they would influence Sally's choices and how she'd be disposed to choose if she were this or that about her choices. (It would be a terrible mistake to suppose that beliefs reduce to or can be analysed as mere choice dispositions, but it would be even more of a mistake to deny that there's a close connection between beliefs and choices, such that we can form reasonably accurate expectations as to how a person will tend to act given assumptions about their beliefs and knowledge of the kinds of things people tend to desire.) And for the large majority of decision situations where the truth or falsity of p is relevant to the outcome, there won't be much difference between believing p to degree 0.34, or 0.37, or anything in-between. So it's entirely plausible that Sally could be in a state where she isn't *quite* how we'd expect if she believed p to degree 0.34, and she isn't *quite* how we'd expect if she believed p to degree 0.37, and so on, but to represent her in any of these ways would still do a good enough job of explaining and predicting her behaviour. They all serve well enough as precisifications of a single vague belief state.

Contrast that with $\mathbf{R}_{\text{wide}}$, where this time $\mathcal{R}(p) = [0, 1]$. Hájek (2003) refers to this as being 'vague over the entire $[0, 1]$ interval'—but what does *that* mean? We have a clear enough idea of what Sally would be like if she was absolutely certain that p is true. And we have a clear enough idea of what Sally would be like if she were absolutely certain that p is false. There isn't much similarity between them, and it's hard to imagine anyone being in a state that comes close enough to playing the *absolutely certain of* p role that $\mu(p) = 1$ does a good job of explaining and predicting her behaviour, whilst also being in a state that comes close enough to playing the *absolutely certain of* $\neg p$ role that $\mu(p) = 0$ also does a good job of explaining and predicting her behaviour. Like the cat and the cassowary, to be in a state that's precisifiable by one is to be in a state that's not precisifiable by the other. And neither of those states is much like being 50% confident in p, or 30% confident, and so on.[4]

Don't say that, where Sally's vague beliefs are given by $\mathbf{R}_{\text{wide}}$, then she'll be in a state that causes her to be *indeterminately disposed* between behaving in the $\mu(p) = 1$ way, and the $\mu(p) = 0$ way, and the $\mu(p) = 0.5$ way, and so on. For what could that possibly mean other than that Sally's *not* disposed to behave in any of those ways at all? Imagine that Sally is considering prices for a dollar bet on p. We could meaningfully say that she's *equally disposed* to accept any price between \$0 and \$1 as fair, or we could say that she *lacks a disposition* one way or the other, but in either case she'll be determinately unlike what we'd expect if she had 0% confidence that p—such an agent wouldn't be willing to

[4] Another analogy, if it helps: it makes sense to say that there's a vague boundary between *being tall* and *not being tall*, but talk of beliefs being 'vague over the entire $[0, 1]$ interval' strikes me as essentially similar to saying that the fuzziness of 'tall' extends all the way from tiniest infants up to the tallest basketballer.

pay anything for the bet. And she'll be determinately unlike what we'd expect if she were 100% confident that p. *All* of the μ in $\mathbf{R}$ misrepresent her belief state, because each has implications regarding the functional role of that state which cannot all jointly be close to the truth. Better to say she's in a *different* belief state that's determinately unlike any of them.

Don't say, either, that vague beliefs are the rational response to vague evidence, and nothing more. It may be true that vague evidence permits only vague beliefs, or it may not be, but either way that's not yet an explanation of what those vague beliefs *are*. If one doesn't know *what it is* for beliefs to be 'vague over the entire [0,1] interval'—or if such a thing even makes sense—then it's not going to help to be told that this is how one's beliefs ought to be when one doesn't have any precise evidence!

(To be clear: there are many potential ways to make sense of extreme indeterminacy of doxastic state. Functionalists will sometimes say that an agent can be in a state that occupies the functional role of pain *for her* even while that very same state occupies the role of pleasure *for her population*, and thus there's simply no fact of the matter as to whether she's really in a state of pain or in a state of pleasure. One could conceive of someone saying the same sort of thing about *believing p to degree 0* and *believing p to degree 1*. Or if you buy into quantum indeterminacy, then no doubt we could construct a *Schrödinger's believer* situation where Sally is in a superposition of radically different belief states. And surely there will be many other ways we could make sense of a representor as representing an indeterminate belief state. But the point here isn't that there's *no* way to make sense of extreme indeterminacy in strength of belief. Rather, the point is that it's unclear how to make sense of extreme indeterminacy in the cases of interest to advocates of the vagueness interpretation—and they're typically interested in indeterminacy as a normal rational response to incomplete or vague evidence, not indeterminacy as a result of some weird quirk of functionalism or quantum-mechanical experiments.)

A good deal of the contemporary discussion around representors involves cases like $\mathbf{R}_{\text{wide}}$, and inasmuch as advocates of the vagueness interpretation want to be a part of that discussion they owe us a sensible explanation of what it is for beliefs to be 'vague over the entire [0,1] interval', and likewise for the interpretation of other representors containing precise functions very far removed from one another. In lieu of a sensible explanation, the alternative is to restrict the application of the vagueness interpretation to those cases where it does make sense—cases like $\mathbf{R}_{\text{narrow}}$—and in so doing, render some of the most interesting literature about imprecise probabilities no longer applicable.

3. The Comparativist Interpretation

Under the vagueness interpretation, it's indeterminate which of the μ in $\mathbf{R}$ is supposed to represent Sally's beliefs—each μ therefore represents Sally's beliefs under a precisification. According to comparativist interpretations, by contrast, if there's more than one μ in $\mathbf{R}$ then each will determinately misrepresent Sally's beliefs. It's the entire set $\mathbf{R}$ which does the representing, and no individual μ within $\mathbf{R}$ has any representational import independent of the whole.

But first I should start with *comparativism*, the idea that absolute degrees of belief are just a way of representing what are ultimately nothing more than relations of relative confidence.[5] To discuss this I'll need some notation:

- $p \succsim q$ iff Sally is at least as confident that p as she is that q
- $p \succ q$ iff Sally is more confident that p than she is that q
- $p \sim q$ iff Sally is just as confident in p as she is in q
- $p \bigtriangledown q$ iff Sally's confidence in p is incomparable to her confidence in q

We assume that $p \bigtriangledown q$ holds whenever p and q are not related by $\succsim$, $\succ$, or $\sim$ in either direction, provided of course they both belong to $\mathbf{P}$. I will thus ignore the possibility that there may be other non-conventional forms of comparability, such as *parity*. I will also take it for granted that if either $p \succ q$ or $p \sim q$, then $p \succsim q$; that seems analytically true if anything is and in any case seems to be common ground among contemporary comparativists. (The other direction is not so obvious.) Given this, $p \not\succsim q$ implies $p \not\succ q$ and $p \not\sim q$, and so it suffices from now on to say

$$p \bigtriangledown q \quad \text{iff} \quad p \not\succsim q \text{ and } q \not\succsim p$$

With that out of the way, we can characterise *traditional comparativism* as the view that a single probability function μ is supposed to represent the facts about Sally's beliefs by virtue of representing her comparative confidences, specifically in the sense that:

$$\begin{Bmatrix} p \succsim q & \text{iff} & \mu(p) \geq \mu(q) \\ p \succ q & \text{iff} & p \succsim q \text{ and } q \not\succsim p \\ p \sim q & \text{iff} & p \succsim q \text{ and } q \succsim p \end{Bmatrix}$$

Note an immediate consequence of traditional comparativism: the $\geq$ relation on the reals is *complete*, in that for any two real numbers a, b, either $a \geq b$ or $b \geq a$; hence, any real-valued function μ automatically represents $\succsim$ as being likewise complete over $\mathbf{P}$. There's no possibility of incomparability on the traditional comparativist picture.

Representors provide us with an alternative means of representing comparative confidence relations, with the benefit of allowing for incompleteness and hence for representing incomparability. Or more accurately: representors provide us with several distinct ways of representing potentially incomplete confidence relations, corresponding to several varieties of comparativist interpretation. One way to capture the difference between them is in terms of which of $\succsim$, $\succ$ and $\sim$ are treated as definitional primitives and which (if any) are treated as derivative. On the traditional single-function model, it's typical to let $\succsim$ be the unique

[5] For discussions sympathetic to comparativism, see (Keynes 1921), (de Finetti 1931), (Koopman 1940b; 1940a), (Fine 1973), (Zynda 2000), and (Stefánsson 2017; 2018). For a recent overview on comparativism and closely related topics, see (Konek 2019). In Konek's terminology, the position being discussed at present is the 'unary measurement-theoretic view'; the 'pluralist measurement-theoretic' version will be discussed at the end of the section. I'd like to enter into the record that there's nothing uniquely *measurement-theoretic* about comparativism; my non-comparativist alternative is at least as firmly grounded in measurement theory as any comparativist interpretation might be.

primitive relation, and simply define $\succ$ and $\sim$ as its asymmetric and symmetric parts respectively (as I did above). But this isn't the only way we could do things. We could just as easily let $\sim$ and $\succ$ be the primitive relations, and define $\succsim$ as the disjunction of the two (i.e., $p \succsim q$ iff $p \succ q$ or $p \sim q$). Or we could treat $\succsim$ and $\succ$ as our primitives and use them to define $\sim$ (e.g., $p \sim q$ iff $p \succsim q$ and $q \succsim p$, or iff $p \not\succ q$ and $q \not\succ p$). Or we could just let all three be considered equally primitive. The point is that *it doesn't matter*—it will make no difference whatsoever which of these we go with when it comes to reading a single function μ as a representation of Sally's comparative confidences. But these choices do make a difference when we shift over to the representor model.

One way to provide a comparativist interpretation of a representor is to treat $\succsim$ as the unique primitive relation. On this interpretation we say that $\mathbf{R}$ represents that $p \succsim q$ just in case every function in $\mathbf{R}$ agrees that p is at least as probable than q, and then we let $\sim$ and $\succ$ be defined out of $\succsim$ as the symmetric and asymmetric parts of $\succsim$ as usual. I'll call this $\succsim$-*comparativism*:

$$\begin{Bmatrix} p \succsim q & \text{iff} & \forall \mu \in \mathbf{R} : \mu(p) \geq \mu(q) \\ p \succ q & \text{iff} & p \succsim q \text{ and } q \not\succsim p \\ p \sim q & \text{iff} & p \succsim q \text{ and } q \succsim p \end{Bmatrix}$$

Consequence: $p \succ q$ just in case $\mu(p) \geq \mu(q)$ for all μ in $\mathbf{R}$, with $\mu(p) > \mu(q)$ for at least one *but not necessarily all* of them. This is probably the most common way of reading a set of probability functions as a representation of comparative probability relations. Or at least it's the way that comes up most often in the literature. See, for example, (Nehring 2009), (Alon & Lehrer 2014), (Miranda & Destercke 2015) (Harrison-Trainor & Holliday 2016), (Bradley 2017), (Harrison-Trainor et al. 2018), (Konek 2019), (Ding et al. 2021), and (Eva & Stern Eva and Stern). We can also find $\succsim$-comparativism implicit in Kaplan's 'Modest Probabilism' (1996; 2002; 2010).[6]

By contrast, Eva (2019, pp. 394-5) puts forward a distinct comparativist interpretation, inconsistent with $\succsim$-comparativism in some cases, according to which $\succ$ and $\sim$ are definitionally primitive and $\succsim$ is just their disjunction. I'll call this one $\succ/\sim$-*comparativism*:

$$\begin{Bmatrix} p \succsim q & \text{iff} & p \succ q \text{ or } p \sim q \\ p \succ q & \text{iff} & \forall \mu \in \mathbf{R} : \mu(p) > \mu(q) \\ p \sim q & \text{iff} & \forall \mu \in \mathbf{R} : \mu(p) = \mu(q) \end{Bmatrix}$$

[6] Kaplan's several slightly different statements of 'Modest Probabilism' all presuppose an interpretation under which (i) $p \sim q$ iff $\forall \mu \in \mathbf{R} : \mu(p) = \mu(q)$, (ii) $p \succ q$ iff $\forall \mu \in \mathbf{R} : \mu(p) \geq \mu(q)$ and $\exists \mu \in \mathbf{R} : \mu(p) > \mu(q)$, and (iii) you are undecided as to the relative credibility of p and q iff $p \not\sim q$, $p \not\succ q$, and $q \not\succ p$. Substituting '$p \triangledown q$' for 'you are undecided as to the relative credibility of p and q,' and assuming as above that ($p \sim q$ or $p \succ q$) implies $p \succsim q$, then Kaplan's (i)–(iii) are just another way of expressing $\succsim$-comparativism.

But wait—there's more! Builes et a. (2022) seem to put forward what we can call $\succsim/\succ$-*comparativism*:[7]

$$\begin{cases} p \succsim q & \text{iff} \quad \forall \mu \in \mathbf{R} : \mu(p) \geq \mu(q) \\ p \succ q & \text{iff} \quad \forall \mu \in \mathbf{R} : \mu(p) > \mu(q) \\ p \sim q & \text{iff} \quad p \succsim q \text{ and } q \succsim p \end{cases}$$

It probably won't be immediately obvious what the impact of these differences will be, but an example will help. Imagine that Sally has been given a coin by a local magician, and has been asked to toss it a few times. She knows that this magician's coins are often biased, though not always, and if it is biased then it will be highly variable in which direction and to what extent. As far as she knows, it it could be completely biased towards heads, or completely biased towards tails, or anything in-between. Given that, we might decide to represent Sally's beliefs by means of a convex representor $\mathbf{R}_{\text{coin}}$ such that, where

p = The coin will land heads on the next toss
q = The coin will land heads on both of the next two tosses

we have $\mathcal{R}(p) = [0, 1]$, and for all μ in $\mathbf{R}_{\text{coin}}$, $\mu(p) = \sqrt{\mu(q)}$. (Don't worry about whether you think this is the *right* way to represent Sally's beliefs in this kind of situation; the important point for the example is that $\mu(p) = \mu(q)$ only where $\mu(p) = 1$ or $\mu(p) = 0$, otherwise $\mu(p) > \mu(q)$.)

Now, since every μ in $\mathbf{R}_{\text{coin}}$ agrees on $\mu(p) \geq \mu(q)$, but they don't all agree on $\mu(q) \geq \mu(p)$, the $\succsim$-comparativist will read $\mathbf{R}_{\text{coin}}$ as saying:

$$p \succsim q, \qquad p \succ q, \qquad p \not\sim q, \qquad p \not\triangledown\ q$$

On the other hand, since neither $\mu(p) > \mu(q)$ nor $\mu(p) = \mu(q)$ for all μ, the $\succ/\sim$-comparativist reads $\mathbf{R}_{\text{coin}}$ as saying that p and q are incomparable:

$$p \not\succsim q, \qquad p \not\succ q, \qquad p \not\sim q, \qquad p \triangledown\ q$$

And on the *other* other hand, since every μ agrees that $\mu(p)$ is at least as much as $\mu(q)$, but not always that it's more than $\mu(q)$, the $\succsim/\succ$-comparativist will read $\mathbf{R}_{\text{coin}}$ as saying:

$$p \succsim q, \qquad p \not\succ q, \qquad p \not\sim q, \qquad p \not\triangledown\ q$$

[7] In more detail: Builes et al. (2022, p. 8) advocate what they call the 'Comparative View', according to which μ belongs to $\mathbf{R}$ iff, (i) if $p \succsim q$ then $\forall \mu \in \mathbf{R} : \mu(p) \geq \mu(q)$, and (ii) if $p \succ q$ then $\forall \mu \in \mathbf{R} : \mu(p) > \mu(q)$. This doesn't yet entail $\succsim/\succ$-comparativism, since without any explicit assumptions about the structure of $\succsim$ and $\succ$ we cannot guarantee the converses of (i) and (ii). However, if we suppose that Sally's comparative confidences are rational in the sense that (a) they can be extended in a way that's representable by at least one probability function, and (b) they do not have any 'gaps' that could be filled by a priori reasoning alone—for instance, if $p \succsim q$ and $q \succsim r$, then it shouldn't be the case that $p \triangledown r$—then the 'Comparative View' will imply the stronger $p \succsim q$ iff $\forall \mu \in \mathbf{R} : \mu(p) \geq \mu(q)$, and $p \succ q$ iff $\forall \mu \in \mathbf{R} : \mu(p) > \mu(q)$. In any case, it's worth noting that clause (ii) of the 'Comparative View' (with or without the rationality assumption) is inconsistent with the more common $\succsim$-comparativism. This plays an important role in the proofs for Builes et al.'s two main theses (pp. 19–20), both of which rely on the premise that $p \succ q$ only if $\mu(p) > \mu(q)$ for *all* μ in $\mathbf{R}$.

The foregoing is useful for highlighting some of the dangers arising from unreflective application of the credal committee metaphor. According to $\succ/\sim$-comparativism and $\succsim/\succ$-comparativism, every voter on the committee needs to agree that $p \succ q$ in order for that to be true, exactly as the metaphor suggests, but this needn't be so for the $\succsim$-comparativist. And if every voter on the committee agrees that $p \succsim q$ then that will be true according to $\succsim$-comparativism and $\succsim/\succ$-comparativism, as the metaphor suggests, but not always according to $\succ/\sim$-comparativism. All three of the comparativist interpretations agree that the credal committee metaphor gets some things right and some things wrong, though they cannot agree on what. And while all three of them agree that $\mathbf{R}$ represents $p \sim q$ only if everyone on the committee votes as such, in contrast to the vagueness interpretation they also imply that $p \sim q$ is (determinately) *false* if so much as a single voter puts their hand up for $p \not\sim q$.

Limitations with the comparativist interpretation

It's uncontroversial that representors can be used to represent incomplete comparative probability orderings. But the comparativist advocates for something stronger: that's *all* a representor represents, and that's all a representor *needs* to represent, because those states of comparative confidence are what ultimately comprise our systems of belief.

One of the most frequently cited motivations for comparativism is the idea that there's nothing about our beliefs that calls for a *unique* numerical representation, or any *numerical* representation at all for that matter (see, e.g., Koopman 1940a, p. 269; Fine 1973, p. 15; Zynda 2000, pp. 64ff; Stefánsson 2017). Builes et al. summarise the thought nicely:

> Comparativism is based on the intuitive thought that while numerical probabilities *represent* belief states, there's nothing about our belief states that mandates a unique numerical representation. In other words, there's nothing "0.69-ish" about my degree of confidence in p, beyond the fact that 0.69 can serve as an adequate representation of my degree of confidence within a particular representational system. But 69, for example, or 732.6 for that matter, would work just as well, provided the system was structured in the right way. (2022, p. 7)

Presumably, when we say that Sally believes p to degree 0.69, we don't thereby imagine that we ought to be able to peek inside her head and find a little '0.69' written somewhere in there—perhaps attached to a sentence in her belief box. Clearly our use of '0.69', or our use of the 0-to-1 scale more generally, is a matter of convention. If we'd settled on a 8.1-to-1050 scale, we'd have said that Sally believes p to degree 732.6. (Maybe that scenario sounds unlikely; but then again, for some reason people still use imperial units.) These are just alternative ways of representing what must be a fundamentally *qualitative* system of beliefs that just so happens to admit of fruitful *quantitative* representation by virtue of some structural isomorphism. So long as the numbers we use appropriately reflect that qualitative structure, it doesn't really matter what scale we use. Comparativism naturally captures this very intuitive thought.

And that *would* be a compelling reason in its favour, if the alternative to comparativism was saying that numerical degrees of belief are not representations but correspond to literal numbers in the head between 0 and 1. But very few people—comparativists and non-comparativists alike—would dispute the idea that our quantitative representations are just a formalised way of capturing what must be a fundamentally qualitative structure characterised by non-numerical properties and relations. None of this is the special province of comparativism. The real debate isn't between those who think the numbers are *representations* and those who think belief systems *literally are* probability functions—that was never at issue! Instead, the interesting debate concerns the nature of the qualitative structures being represented. What exactly are the non-numerical properties and relations captured by our numerical representations of belief, and what, therefore, are the properties of those representations which must be shared among any representational alternatives with an equal claim to adequacy?

Comparativists put forward a specific account of what that qualitative structure must be—it is an ordering over a space of propositions imposed by those states of comparative confidence that ultimately make up our beliefs. What matters is that the representation gets the ordering right. Comparativists sometimes also say that numerical ratios and/or intervals of ratios in the numbers assigned ought to reflect how the confidence ordering interacts with unions of disjoint propositions (see Stefánsson 2017, pp. 4–5; Elliott 2022a, pp. 2847–50), but even in this case what matters is the ordering. And that is a very strong claim indeed. Why should we suppose that the full extent of the relevant *qualitative* structure being represented by a probability function, or by a representor, can be expressed wholly in terms of comparative confidence relations?

This question becomes especially pressing once we recognise that contemporary theories of rational decision-making—both for precise and imprecise degrees of belief—make regular appeal to properties of numerical representations that cannot be reduced to or expressed in terms of mere relations of comparative confidence. Here's an example that's as simple as it gets. (It's not the only example, nor the most compelling, but it's easy and gets the point across.) Let **P** be the smallest non-trivial algebra it can be, consisting of four propositions: p, its (relative) complement $\overline{p}$, the contradiction $p \cap \overline{p}$, and the tautology $p \cup \overline{p}$. We suppose that

$$(p \cup \overline{p}) \ \succ \ p \ \succ \ \overline{p} \ \succ \ (p \cap \overline{p})$$

According to all three varieties of comparativism I've discussed, a representor **R** will determine these comparative confidences if (but not in all cases only if), for all μ in **R**, $1 > \mu(p) > \mu(\overline{p}) > 0$. There are infinitely many probability functions that will represent those comparative confidences, and many more representors that will do the same—including but not limited to any set such that $\mathcal{R}(p)$ is a sub-interval of $(0.5, 1)$. So they all represent the very same system of beliefs, according to comparativism.[8]

[8] There will be only one *maximally inclusive* **R** that represents this ordering, and comparativists may want to say that we should always use the maximally inclusive representor. This makes no difference to my argument. The point also does not depend on **P** being finite. For any

But now take any decision theory for imprecise probabilities that generalises expected utility theory—specifically in the sense that it includes expected utility theory as a special case when $\mathbf{R}$ is a singleton set. That theory will entail that there is a decision-theoretically relevant difference between at least some of these representors. For instance, imagine you have a choice between two gambles. On the first gamble, you get \$1 if p is true and nothing otherwise. On the second, you get \$2 if p is false, and nothing otherwise. Which should you choose? Case 1: if $\mathcal{R}(p)$ picks out an interval in the upper end of the $(0.5, 1)$ range—say, $\mathcal{R}(p) = [0.98, 0.99]$—then you should prefer the first gamble. Case 2: $\mathcal{R}(p)$ picks out an interval in the lower end of that range—say, $\mathcal{R}(p) = [0.51, 0.52]$—then you should prefer the second gamble. Case 3: if $\mathcal{R}(p)$ picks out the entire $(0.5, 1)$ interval, then depending on the theory of decision making with imprecise probabilities you choose to go with you might either be indifferent between the two gambles, prefer the first to the second, the second to the first, or lack a preference—but either way you'll get a different result from Case 1 or Case 2, or both.

I said that there will be a decision-theoretically relevant difference between *at least some* representors, for *any* decision theory of the specified kind. The proof is simple. Expected utility theory recognises distinctions in ordinally-equivalent probability functions—preferences licensed by the function μ when conjoined with a utility function need not be licensed by μ' when conjoined with those same utilities, even if μ and μ' determine the same ordering. So, any theory that generalises expected utility theory appeals to properties of the singleton sets of μ and μ' that cannot be reduced to the identical comparative confidence orderings they determine. QED. Those theories make appeal to what I'll call *extra-ordinal* information. But the conclusion of that proof is also misleadingly weak. For any of the well-known accounts of decision-making with imprecise probabilities, such as Γ-maximin, E-admissibility, maximality or interval dominance—see (Troffaes 2007) for an overview—we can replace the 'at least some' with 'many', and in some cases with 'all'. In particular, pairs of *non*-singleton representors $\mathbf{R}$ and $\mathbf{R}'$ which determine the same comparative confidence orderings can *and often do* have distinct roles to play within these theories.

So the role representors play in those theories requires them to carry meaningful information going beyond what can be expressed merely in terms of comparative confidence relations. By denying the meaningfulness of that additional information, comparativism leaves us with an impoverished picture of what our doxastic states are really like, and one that's too thin to play the theoretical roles seemingly required of it. An *adequate* representation of our beliefs is one that appropriately captures *all* of the qualitative properties and relations upon which the theoretical roles of our beliefs depend, and if we take our current theories of decision-making seriously then the theoretical roles of our beliefs in relation to our desires in the production of intelligent action requires us to suppose they have *some* meaningful qualitative structure extending beyond just the orderings they determine.

$\mathbf{P}$, countable or uncountable, there will always be complete confidence orderings that can be represented by many distinct probability functions, and incomplete confidence orderings that can be represented by many distinct representors.

"Not so fast", you might be thinking to yourself, "Aren't you just presupposing that these theories of decision-making are *correct* in making appeal to this additional information, and therefore begging the question against comparativism? And don't we have some reason already to suppose that those theories often help themselves to more information than they're entitled, as for instance when they represent decision-makers as having complete awareness of their state-space, or when they employ infinitely precise distinctions in degree of belief? Decision theory is rife with idealisations—so what's to stop the comparativist from simply saying that, inasmuch as one of these theories does make use of extra-ordinal information, then that is yet another idealisation?"[9]

In response, we can either focus on the forest or on the trees. Let's start with the trees. While it's clearly true that current theories of decision-making are unrealistic and over-idealising in many respects *when applied to ordinary agents*—as when they presume full awareness, for instance—those same idealisations still seem to have a perfectly sensible interpretation in conceivable scenarios involving idealised agents. While we might not have full awareness of our state-space, an appropriately idealised agent presumably could. And while we probably don't have infinitely precise gradations in degree of belief, there's no reason to think such is impossible. The 'extra information' in these cases isn't *meaningless*—it's not a mere artefact of the model, but something that has a legitimate role to play in some conceivable applications. The situation with comparativism is quite different. Comparativists aren't saying that there's further information encoded in a probability function, or in a representor, which makes sense for idealised agents but not for us. They're saying that there's nothing more to a system of beliefs than the comparative confidences comprising it. If a decision theory treats μ and μ' differently, or $\mathbf{R}$ and $\mathbf{R}'$ differently, despite those fixing the same comparative confidence orderings, then that isn't the theory being inapplicable to mere mortals—it is the theory making appeal to information that doesn't have an interpretation in *any* scenario no matter how idealised. I'll happily accept that contemporary theories of decision-making are often unrealistic when applied to ordinary agents, I don't yet see a compelling reason to suppose they're fundamentally unsound in application to idealised decision-makers.

Now the forest: yes, you could say I'm presupposing that our current decision theories are correct in appealing to that additional information, and indeed there's nothing in principle stopping a comparativist from simply denying that this is legitimate. No questions are being begged, however, because the point of the argument isn't to convince you that comparativism cannot be true; rather, it's to establish that we have good reasons for thinking our numerical representations of belief carry meaningful information beyond that which can be ex-

[9] I am paraphrasing a referee's comments here. The referee suggested, in light of those comments, that I should simply omit the entire line of argument relating to decision theory and instead "stick to the argument that there are doxastic states that cannot be expressed in comparativist terms". Consequently, I feel it's worth being explicit that my argument that there's more to those doxastic states than can be expressible in terms of comparative confidence relations *just is* the argument that the role played by beliefs in our current theories of rational decision-making requires them to possess more structure than comparativism allows. There might be other arguments, but as a functionalist *that* is the one I find compelling.

pressed in terms of comparative confidence relations. Current theories of rational decision-making—not just one or a few of them, but the majority—imply that there can be relevant differences between ordinally-equivalent probability functions and/or representors. These theories thus implicitly characterise distinct but ordinally-equivalent systems of belief, by virtue of the distinctive roles played by the ordinally-equivalent representations therein. Until something better comes along, we have every reason to take that fact seriously as telling us something important about what the numbers represent.

Of course it's *possible* that contemporary theorists have been misled by the additional structure of encoded in our standard numerical representations, and there really aren't any legitimate differences between belief states that don't ultimately reduce to differences in comparative confidence orderings. And *maybe* there's a better theory that comparativists will one day develop—one that never requires going beyond ordinal information, and does at least as good a job of explaining and predicting rational decision-making, of fitting with intuitions and exemplifying the various theoretical virtues, as anything we've got now. Maybe. Possibly. One day. In the meantime, I'm going to take our current theories at face value—and do so with good reason—rather than resting my metaphysics of mind on unsubstantiated hopes and dreams.[10]

A final note: I have focused on the role of representors in decision theory, but the same applies to the theoretical role of probability functions and representors in much of contemporary epistemology. Example one: the relation of probabilistic independence is important for our theories of evidence and learning, but cannot be fully captured in terms of binary comparative confidence relations. (See Joyce 2010, pp. 285ff, for a discussion on this point and several other species of doxastic attitudes and relations that cannot be properly formulated in terms of binary comparative confidence relations.) Example two: epistemic utility theory appeals to properties that differentiate ordinally-equivalent probability functions, and there are at yet no plausible epistemic arguments for imposing constraints on comparative confidences that suffice to guarantee a unique probabilistic representation. (See Mayo-Wilson & Wheeler 2019, p. 19, for discussion.) Example three: as there are meaningfully distinct but ordinally-equivalent objective chance functions, the Principal Principle (Lewis 1980) presupposes distinctions in rational doxastic states that make no sense under comparativism. Example four: theories of peer disagreement require interpersonal comparability of beliefs, and thus far comparativists have no plausible theory of interpersonal confidence comparisons. (See Elliott 2022b for discussion.) And so on, and on, and on.

None of this in any way implies that there's only one way to numerically model our beliefs, or that we must posit fundamentally numerical properties and relations to characterise our beliefs. We can all agree that what's *real* in a numerical representation are the qualitative structures it represents, whatever those structures may be. All it means is that we shouldn't suppose that there's nothing more to our beliefs than relations of comparative confidence.

[10] Fine (1973, pp. 37ff) goes some of the way towards developing a fully comparativist theory of decision-making—though as he notes, 'clearly much remains to be done' (1973, p. 16). That is as true today as it was back then.

The pluralist retreat

One response to these sorts of worries is to back down from the strong thesis that comparative confidences are the uniquely fundamental kind of doxastic state, and instead claim that they are merely one species of fundamental doxastic state among others. These others may include, e.g., qualitative judgements of probabilistic independence, or judgements that one proposition is evidence for another, or non-comparative states of certainty or full belief and disbelief, and so on. Joyce (2010) seems to suggest a position like this. Konek (2019, pp. 308ff) refers to it as the *pluralist* view; Elliott (2022a) calls it *supplemented comparativism*. Whatever we call it, though, it trivially avoids the problems above—if comparative confidences are insufficient, just add more kinds of doxastic state until you've got enough!

But that doesn't mean we ought to be pluralists. Distinguish:

Representational Pluralism. Our numerical models of belief serve to represent multiple kinds of doxastic state, including at least comparative confidences, but also other forms of doxastic attitude that cannot be reduced to comparative confidences.

Fundamentality Pluralism. Our numerical models of belief serve to represent multiple kinds of *fundamental* doxastic state, including at least comparative confidences but also other attitudes that cannot be reduced to comparative confidences.

The foregoing arguments support representational pluralism over its negation, and they support fundamentality pluralism over non-pluralist comparativism. But what they don't do is support fundamentality pluralism over the much weaker representational pluralism, because they don't provide any reasons for thinking that comparative confidences are a species of *fundamental* doxastic state. I think that *everyone* should accept representational pluralism. Fundamentality pluralism is another matter.

The closest that I'm aware of to an actual *argument* for the fundamentality of comparative confidence goes as follows.[11] Say that an *absolute degree of belief* is an attitude connecting an agent to a single proposition and a degree or strength, which may or may not be expressed numerically. These are the kinds of attitudes we attribute when we say, for example, that Sally is certain that p, or very confident that p, or believes p to degree 0.69. Given that, the first premise: a system of absolute degrees of belief always determines a corresponding system of comparative confidences—e.g., if Sally's degree of belief in p is greater than her degree of belief in q, then $p \succ q$. Second premise: the very idea of *degrees* of belief

[11] Fine briefly mentions something like this: '(2) [Comparative probability] provides a wider class of models of random phenomena than does the usual quantitative theory. [...] Point (2) refers to the curious phenomenon that there exist relatively simple examples of what we consider to be valid [comparative probability] statements that are incompatible with any representation in the usual quantitative theory' (1973, pp. 15–6). As stated the 'wider' claim not quite correct. Since there are distinct probability functions/representors that determine identical comparative probability relations, the class of purely relational models cannot be wider than the class of precise or imprecise numerical models, just different.

presupposes that those degrees have a minimal relational structure—they must constitute at least a preorder, so transitivity and reflexivity are required. Anything less and we'd be stretching the usual notion of *degrees* beyond recognition. Third premise: non-transitive systems of comparative confidence are conceivable. So, while every conceivable system of absolute degrees of belief corresponds to a system of comparative confidences, the same is not true in reverse; hence, it cannot be the case that the facts about comparative confidence reduce to the facts about absolute degrees of belief.

That's a really good reason to suppose that absolute degrees of belief are not more fundamental than comparative confidences. But it isn't yet enough to conclude that comparative confidences are more fundamental than absolute degrees of belief—still less that they are fundamental *simpliciter*. For all I've said, comparative confidences and absolute degrees of belief both may be reducible to yet another kind of doxastic state that's still more fundamental than either of them. For example, Suppes & Zanotti (1976) show how one might potentially derive both absolute degrees of belief and comparative confidences out of appropriately structured qualitative expectations. Or one might suppose that both absolute degrees of belief and comparative confidences fall out of the facts about our outright beliefs under special circumstances (cf. Easwaran 2016). Or maybe, as Lewis often suggested (e.g., 1986, pp. 36–7; 1994, p. 430), we can see the system of beliefs *as a whole* as comprising the fundamental doxastic unit. On this picture, which I'll be advocating shortly, our talk of comparative confidences, absolute degrees of belief, outright beliefs and so on are all ultimately just ways of describing salient aspects of a single rich belief state characterised by its functional role in relation to evidence and intentional behaviour.

No doubt there are other possibilities still to be imagined. For the present point it doesn't matter which is correct; what matters is that they're all still on the table. And inasmuch as that's the case, the claim that comparative confidence must a fundamental species of doxastic attitude, whether uniquely so or not, is a dogma of non-pluralist comparativism and fundamentality pluralism alike.

4. The Functional Interpretation

Numerical models of belief are representations of *some* qualitative structure relating to our beliefs, I've said. There's two ways we could think of that structure. On the one hand, it might be something to do with the intrinsic structure of the total system of beliefs considered in isolation. When comparativists say that probability functions and representors represent complete and incomplete comparative confidence orderings respectively, they're focusing on intrinsic structure in this sense. When fundamentality pluralists say that our numerical models also represent other doxastic states not reducible to comparative confidence, they too are focusing on intrinsic structure. A rather different approach, though, is to suppose that the models represent not the intrinsic structure of the belief system, or not *only* that, but also something about how our belief systems relate to other stuff—desires, preferences, evidence, for example. Sometimes it's not possible to understand what a numerical model of some phenomenon represents without understanding how that phenomenon interacts with others.

This idea is not new. The theory of conjoint measurement was designed to show how interactions between quantities can give rise to meaningful information that's not apparent when we consider each quantity in isolation (Debreu 1960; Luce & Tukey 1964). Consider first the familiar case of length. Lengths are measured on a ratio scale, and thus transformations between normal measures of length (feet, miles, meters, parsecs, etc.) will always preserve ratios. This is not idle stipulation: the *twice the distance* relation has genuine meaning, which we can appreciate simply by considering how lengths relate to one another. If Spot the dog is twice as long as Fluffy the cat, then if we were to have two copies of Fluffy and line them up them head-to-tail, their combined length would be exactly as long as Spot. In that sense, one need only look at the intrinsic structure of the system of lengths-relations to appreciate why ratios of lengths have real-world meaning. If comparativism is right, then the measurement of confidence is essentially similar to the measurement of length.

But things need not always work that way. Assume that two quantities $\mathbf{A}$ and $\mathbf{B}$ lack the kind of intrinsic structure had by the system of length-relations that would usually be used to justify their measurement on a ratio scale. Still we might consider how $\mathbf{A}$ and $\mathbf{B}$ trade-off to produce varying levels in a third quantity, $\mathbf{C}$. Let a_1, a_2, a_3 be distinct levels of $\mathbf{A}$, b_1, b_2, b_3 be distinct levels of $\mathbf{B}$, and a_ib_j be the level of $\mathbf{C}$ produced by the interaction of a_i and b_j. Suppose, for some a_1b_1 less than a_2b_1, that $a_1b_2 = a_2b_1$. So the shift from a_1 to a_2 produces the same effect in $\mathbf{C}$ as the shift from b_1 to b_2. Provided all three quantities are qualitatively related to one another in the right kind of way, we can say the difference between a_2b_2 and a_1b_1 is *twice* the difference between a_1b_2 and a_1b_1, or between a_2b_1 and a_1b_1. Now if there are a_3 and b_3 such that $a_1b_3 = a_2b_2 = a_3b_1$, then the shift from a_1 to a_3 produces twice the difference as the shift from a_1 to a_2; likewise for b_1, b_2 and b_3, *mutatis mutandis*. Thus we can show that ratios in the numbers used to represent the different levels of $\mathbf{A}$ and $\mathbf{B}$ can have just as much genuine meaning as ratios in the measurement of length, but that meaning will be apparent only when observing how those quantities relate to one another and to $\mathbf{C}$. If you want to know what the numbers mean when measuring $\mathbf{A}$, or $\mathbf{B}$, then you cannot just consider $\mathbf{A}$ or $\mathbf{B}$ in isolation.[12]

In light of that, consider a well-known argument from Lyle Zynda (2000). Zynda observes that whenever an agent is representable as making decisions in line with the hypothesis that she's an expected utility maximiser with a system of beliefs given by some probability function μ, we can *also* represent her as making decisions in line with some other decision rule given a system of beliefs represented by some alternative and non-probabilistic function. That is, suppose first that Sally's choices fit the hypothesis that she's an expected utility maximiser with beliefs represented by the probability function μ and desires represented by the utility function u. Now define the *believability ranking* β, such that for all p in $\mathbf{P}$, $\beta(p) = 9\mu(p) + 1$. The function β represents degrees of belief on a scale from

[12] One of the best-known applications of the theory of conjoint measurement is in decision theory: Kahneman & Tversky (1979) use it to develop the axiomatic basis for prospect theory, showing how the decision weights and utility functions in prospect theory can be measured on ratio scales by virtue of their joint interaction in the production of preferences.

1 to 10, with $\beta(p \cup q) \neq \beta(p) + \beta(q)$ even when p and q are mutually exclusive. If β is combined with u according to the normal expected utility rule, then it will lead to *very* different preferences than what we get if we combine μ and u according to that same rule. However, say that Sally is a *valuation maximiser* iff, where her beliefs and desires are given by β and u respectively, and for acts a and a' which have outcomes o_i and o'_i respectively at states s_i, Sally weakly prefers a to a' just in case

$$\sum_i \Big[\beta(s_i)u(o_i) - u(o_i)\Big] \geq \sum_i \Big[\beta(s_i)u(o'_i) - u(o'_i)\Big]$$

And if Sally's choices fit the hypothesis that she's an expected utility maximiser with beliefs and desires given by μ and u, then they will also fit the hypothesis that she's a valuation maximiser with beliefs and desires given by β and u. So you can make changes to way beliefs are represented without affecting the preferences they generate, provided you make corresponding adjustments to how the decision rule is formulated. (See also Joyce 2015, p. 419, for similar points.)

The lesson that Zynda draws from this observation is that neither μ nor β are the uniquely correct representation of Sally's beliefs, and that what's *real* are the qualitative properties common to them and any similar 'redefinitions'. In particular, since β is a positive affine transformation of μ, β and μ will impose the very same comparative confidence ordering over $\mathbf{P}$. Thus,

> According to this solution, people really have properties that can properly be called 'degrees of belief', though these are more abstract in nature than subjective [numerical] probabilities, being purely qualitative … *The concept of degree of belief on this strategy becomes a purely ordinal notion* (although it remains the case that rational degrees of belief would always have a cardinal representation.) (2000, p. 65, emphasis added)

But there was a leap there. While it's true that a probability function and the corresponding 'believability ranking' will always share their ordinal structure, that's not *all* they share. The affine transformations that link probability functions and 'believability rankings' preserve lots of properties, not just the orderings. Most importantly, the transformation from μ to β is bijective, so $\mu(p) \neq \mu(q)$ iff $\beta(p) \neq \beta(q)$ and consequently if μ and μ' are any two non-identical probability functions, then their corresponding 'believability rankings' will likewise be non-identical: $\mu \neq \mu'$ iff $\beta \neq \beta'$. This is true even if μ and μ' are ordinally equivalent. And in just the same way that differences between μ and μ' can matter for some decision situations when using the expected utility rule, the differences between β and β' can matter for some decision situations when using the valuation maximisation rule. So Zynda's argument *doesn't* support treating the concept of *degree of belief* as 'a purely ordinal notion' after all.

The reader may note that $\mu(p) \neq \mu(q)$ iff $\beta(p) \neq \beta(q)$ precisely because affine transformations preserve ratios of differences. That is, for any real value r,

$$\mu(p) - \mu(q) = r\big[\mu(x) - \mu(y)\big] \text{ iff } \beta(p) - \beta(q) = r\big[\beta(x) - \beta(y)\big]$$

But do not place any weight on this fact—if Zynda's example shows that μ and β are equally good representations of Sally's beliefs, then a parallel example shows that a representation of Sally's beliefs needn't be an affine transformation of μ at all. Suppose $\gamma(p) = \beta(p)^2$. If Sally can be represented as an expected utility maximiser with beliefs given by μ, or as a valuation maximiser with beliefs given by β, then she can also be represented as having beliefs given by the function γ and following the decision rule which tells her to weakly prefer a to a' iff

$$\sum_i \left[\sqrt{\gamma(s_i)} - 1\right] u(o_i) \;\geq\; \sum_i \left[\sqrt{\gamma(s_i)} - 1\right] u(o_i')$$

Ratios of differences are not shared between μ and γ. Nor, for that matter, does the example even require the alternative representations preserve the ordering of the probabilistic representation! It suffices that the transformation is a bijective automorphism, and so invertible. For any transformation which takes us from a probability function μ to some alternative δ, if it's such that $\mu(p) \neq \mu(q)$ iff $\delta(p) \neq \delta(q)$, then there will be at least one (potentially very complicated) decision rule which generates the same preferences when using the 'redefined' model δ as the expected utility rule does when using μ.

The lesson of Zynda's example isn't that what's *real* in the systems of belief represented by μ and β and γ (and δ and …) are the qualitative properties that are shared between all of them. There are basically *zero* qualitative properties common across all of them, aside from the utterly trivial requirement that different degrees of belief will be assigned different numerical values. Rather, the lesson should have been that when we are modelling beliefs in decision theory, the structure we're trying to represent is not something internal to system of beliefs itself, considered in isolation from anything else, but instead at least partly something about the relations that hold between beliefs, desires, and preferences. *That* is why we cannot make alterations to the probabilistic model of beliefs without making corresponding adjustments to the decision rule—because the *meanings* of the probabilities in the model are tied up with how they interact with the utilities to produce preferences. What's common to the probabilistic μ and the non-probabilistic β and γ are the similar roles they play in the models of decision-making that make use of them respectively—they interact with *these* utilities to produce *those* preferences. What the numbers need to represent is that functional role; how they represent it is up to us.

So here's my thought: if one cannot appreciate what is *real* versus what is a mere *artefact* in a formal model of belief without appreciating the role those models play in the theories that make use of them, then why not just take those roles themselves to be what's real? We don't *have* to come up with an interpretation of representors that's independent of the theories in which they figure—an interpretation cashed out in terms of vagueness or comparative confidence or anything else—since the interpretation of $\mathbf{R}$ can just be *the system of beliefs that plays the* $\mathbf{R}$*-role in that theory.* Thus, I claim: if (and only if) our best theories of rational belief and decision-making posit distinctive roles for $\mathbf{R}$ and $\mathbf{R}'$, then the differences between $\mathbf{R}$ and $\mathbf{R}'$ mark a genuine difference in the doxastic states they represent.

Sometimes the differences in the systems of belief represented by $\mathbf{R}$ and $\mathbf{R}'$ might be expressible as differences of intrinsic structure—such as, for example, differences in the comparative confidences they determine. But we do not have to suppose that differences between $\mathbf{R}$ and $\mathbf{R}'$ always reflect differences in intrinsic structure. Sometimes the best we might be able to say is that $\mathbf{R}$ and $\mathbf{R}'$ just represent different systems of belief, as evidenced by their distinct roles in the production of preferences for example, even if there's no simple way of describing those differences besides describing the divergent parts they play in rational decision-making. Still less do we have to assume that meaningful differences between $\mathbf{R}$ and $\mathbf{R}'$ will be expressible in ordinary language.

For example, consider $\mathbf{R}_{coin}$ once more. Should we say $\mathbf{R}_{coin}$ represents (i) that Sally is more confident in p than she is in q, as the $\succsim$-comparativist suggests; (ii) that she's at least as confident in p as she is in q but neither strictly more nor equally so, as the $\succsim/\succ$-comparativist suggests; or (iii) that p and q are simply incomparable, as the $\succ/\sim$-comparativist suggests? I say there's no fact of the matter. Our intuitive conception of gradability—of one thing's being *more than* or *at least as much as* another—is built atop the model of $>$ and $\geq$, which presupposes a complete ordering. Thus we tend to presuppose that $p \succ q$ iff $p \succsim q$ and $q \not\succsim p$, that $p \succsim q$ iff $p \succ q$ or $p \sim q$, that $p \not\succsim q$ iff $q \succ p$, and so on. But when dealing with incomplete orderings the usual presuppositions break down, and we're left with several ways of defining $\succ$ and $\succsim$ in relation to one another—all of them preserving some aspects of the ordinary conception while dropping others, giving rise to counterintuitive implications (e.g., that $p \succsim q$ can be true even though both $p \succ q$ and $p \sim q$ are false). The alternative definitions are all roughly as good (and as bad) as each other, so I say it's indeterminate how we should describe Sally's comparative confidences if her beliefs are represented by $\mathbf{R}_{coin}$. This doesn't mean that Sally's *beliefs* are indeterminate, though. The facts about her beliefs are settled by saying they're represented by $\mathbf{R}_{coin}$, and are therefore the kinds of belief that play the $\mathbf{R}_{coin}$-role. What's unsettled is how we ought to use phrases like 'more confident' in the description of those beliefs when her comparative confidence relation is incomplete.

Eriksson & Hájek (2007, pp. 204ff) once proposed something like what I have in mind here. What they propose is that (absolute) degrees of belief are those things that play the kinds of roles numerical probabilities are supposed to play in the best systematisations of our ideas about rational belief and decision-making. They called their view *primitivism*, but they also note (2007, p. 210) the idea is very much in the spirit of the Canberra plan, or functionalism more generally—the main difference being that the functionalist will want to say that our best theories implicitly define what 'degrees of belief' are via their distinctive roles, whereas they question whether these 'definitions' should really be counted as such. They prefer instead to say that the concept of 'degrees of belief' is a theoretical primitive, and we get a handle on the concept by understanding the roles it plays in the theories that make use of them. It is a difference that makes little difference. The essence of Eriksson & Hájek's proposal is functionalism, broadly construed, and in that respect is closely related to ours.

So call my position *primitivism* if you will. But what I'm proposing is not quite what Eriksson & Hájek suggested. In their case, the proposed primitives are *absolute degrees of belief.* That makes sense inasmuch as we're modelling beliefs in the traditional way, since everything a probability function says about a belief state can be derived from what it says about the absolute degree of belief it associates with each proposition. But if our best theories instead make use of representors, then we'd be wise not to take individual degrees of belief as our 'theoretical primitives'. What a representor represents cannot always be captured merely by specifying what it says about the (imprecise) degree to which the agent believes each proposition. That is what the summary function $\mathcal{R}$ does, but a summary function can omit information relevant to how belief state is structured as a whole and the role it plays. $\mathbf{R}_{coin}$ assigns the very same maximally imprecise interval to p and q, but it would be a mistake to say that Sally's attitudes towards p and q are the same. Better instead to follow Lewis: let the entire system of beliefs be our primitive, represented by the set of functions $\mathbf{R}$, and characterise that total system of beliefs by the functional role played by $\mathbf{R}$ in the best theories we have that make use of such models.

And what, finally, are the 'best' theories to which I keep referring? That's for you to decide. What I'm offering is a schema for an interpretation—of the many and various theories which make use of representors, pick the ones you have the most reasons to like and the interpretation of $\mathbf{R}$ will fall out of the role that $\mathbf{R}$ plays therein. If it turns out that a comparativist theory of rational belief and decision-making is best, then the functional interpretation should just reduce to the comparativist interpretation. But I think that's unlikely. Too much of contemporary decision theory and formal epistemology draws on the extra-ordinal properties of our numerical representations for me to confidently suppose otherwise. For much the same reason, I think it likely that our best theories will make use of representors (like $\mathbf{R}_{wide}$ and $\mathbf{R}_{coin}$) which posit 'highly imprecise' degrees of belief. And if that's true, then the vagueness and comparativist interpretations are not up to the task.

References

Alon, S. and E. Lehrer (2014). Subjective multi-prior probability: A representation of a partial likelihood relation. *Journal of Economic Theory 151*, 476–492.

Bradley, S. (2014). Imprecise Probabilities. *Stanford Encyclopedia of Philosophy*.

Bradley, Richard (2017). *Decision Theory with a Human Face.* Cambridge University Press.

Builes, D., S. Horowitz, and M. Schoenfield (2022). Dilating and contracting arbitrarily. *Nous 56*(1), 3–20.

Chang, R. (2002). The Possibility of Parity. *Ethics 112*(4), 659–688.

de Finetti, B. (1931). Sul Significato Soggettivo Della Probabilita. *Fundamenta Mathematicae 17*(1), 298–329.

Debreu, G. (1960). Topological methods in cardinal utility theory. In K. J. Arrow, S. Karlin, and P. Suppes (Eds.), *Mathematical Methods in the Social Sciences*, pp. 16–26. Stanford University Press.

Ding, Y., W. H. Holliday, and T. F. Icard III (2021). Logics of Imprecise Comparative Probability. *International Journal of Approximate Reasoning 132*, 154–180.

Easwaran, K. (2016). Dr. Truthlove, or: How I Learned to Stop Worrying and Love Bayesian Probabilities. *Nous 50*(4), 816–853.

Elliott, E. (2022a). Comparativism and the Measurement of Partial Belief. *Erkenntnis 87*, 2843–2870.

Elliott, E. (2022b). What is 'Real' in Interpersonal Comparisons of Confidence. *Australasian Journal of Philosophy 100*(1), 102–116.

Eriksson, L. and A. Hájek (2007). What are degrees of belief? *Studia Logica 86*(2), 183–213.

Eva, B. (2019). Principles of Indifference. *The Journal of Philosophy 116*(7), 390–411.

Eva, B. and R. Stern. Comparative Opinion Loss. *Philosophy and Phenomenological Research*.

Fine, T. L. (1973). *Theories of Probability: An Examination of Foundations*. Academic Press.

Fishburn, P. C. (1986). The axioms of subjective probability. *Statistical Science 1*(3), 335–345.

Harrison-Trainor, M., W. H. Holliday, and T. F. Icard (2016). A Note on Cancellation Axioms for Comparative Probability. *Theory and Decision 80*(1).

Harrison-Trainor, M., W. H. Holliday, and T. F. Icard III (2018). Inferring Probability Comparisons. *Mathematical Social Sciences 91*, 62–70.

Hájek, A. (2003). What Conditional Probability Could Not Be. *Synthese 137*(3), 273–323.

Hájek, A. and M. Smithson (2012). Rationality and indeterminate probabilities. *Synthese 187*, 33–48.

Jeffrey, R. (1983). Bayesianism with a human face. *Testing Scientific Theories, Minnesota Studies in the Philosophy of Science 10*, 133–56.

Joyce, J. (2005). How Probabilities Reflect Evidence. *Philosophical Perspectives 19*, 153–178.

Joyce, J. (2010). A Defense of Imprecise Credences in Inference and Decision Making. *Philosophical Perspectives 24*(1), 281–323.

Joyce, J. (2015). The Value of Truth: A Reply to Howson. *Analysis 75*(3), 413–424.

Kahneman, D. and A. Tversky (1979). Prospect theory: An analysis of decision under risk. *Econometrica 47*(2), 263–291.

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

Kaplan, M. (2002). Decision theory and epistemology. In P. Moser (Ed.), *The Oxford Handbook of Epistemology*, pp. 434–462. Oxford: Oxford University Press.

Kaplan, M. (2010). In defense of modest probabilism. *Synthese 176*(1), 41–55.

Keynes, J. M. (1921). *A Treatise on Probability.* New York: Macmillan.

Konek, J. (2019). Comparative Probabilities. In *The Open Handbook of Formal Epistemology*, pp. 267–348. PhilPapers Foundation.

Koopman, B. O. (1940a). The Axioms and Algebra of Intuitive Probability. *Annals of Mathematics 41*(2), 269–292.

Koopman, B. O. (1940b). The Bases of Probability. *Bulletin of the American Mathematical Society 46*(10), 763–774.

Levinstein, B. (2019). Imprecise Epistemic Values and Imprecise Credences. *Australasian Journal of Philosophy 97*(4), 741–760.

Lewis, D. (1980). A Subjectivist's Guide to Objective Chance. In R. Jeffrey (Ed.), *Studies in Inductive Logic and Probability*, pp. 263–293. University of California Press.

Lewis, D. (1986). *On the Plurality of Worlds.* Cambridge University Press.

Lewis, D. (1994). Reduction of Mind. In S. Guttenplan (Ed.), *Companion to the Philosophy of Mind*, pp. 412–431. Blackwell.

Lewis, D. (1999). Many, but almost one. In *Papers in Metaphysics and Epistemology*, pp. 164–182. Cambridge University Press.

Luce, R. D. and J. W. Tukey (1964). Simultaneous conjoint measurement: a new scale type of fundamental measurement. *Journal of Mathematical Psychology 1*(1), 1–27.

Mayo-Wilson, C. and G. Wheeler (2019). Epistemic Decision Theory's Reckoning. Manuscript. http://philsci-archive.pitt.edu/16374/1/25a_EDTR.pdf.

Miranda, E. and S. Destercke (2015). Extreme points of the credal sets generated by comparative probabilities. *Journal of Mathematical Psychology 64-65*, 44–57.

Nehring, K. (2009). Imprecise probabilistic beliefs as a context for decision-making under ambiguity. *Journal of Economic Theory 144*, 1054–1091.

Rinard, S. (2015). A Decision Theory for Imprecise Credences. *Philosopher's Imprint 15*(7), 1–16.

Rinard, S. (2017). Imprecise Probability and Higher Order Vagueness. *Res Philosophica 94*(2), 257–273.

Seidenfeld, T. (1988). Decision Theory Without 'Independence' or Without 'Ordering'. *Economics and Philosophy 4*(2), 267–290.

Stefánsson, H. O. (2017). What Is 'Real' in Probabilism? *Australasian Journal of Philosophy 97*(3), 573–587.

Stefánsson, H. O. (2018). On the Ratio Challenge for Comparativism. *Australasian Journal of Philosophy 96*(2), 380–390.

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

Suppes, P. and M. Zanotti (1976). Necessary and sufficient conditions for existence of a unique measure strictly agreeing with a qualitative probability ordering. *Journal of Philosophical Logic 5*(3), 431–438.

Troffaes, M. C. M. (2007). Decision making under uncertainty using imprecise probabilities. *International Journal of Approximate Reasoning 45*, 17–29.

van Fraassen, B. (1990). Figures in a Probability Landscape. In J. Dunn and A. Gupta (Eds.), *Truth or Consequences*, pp. 345–356. Amsterdam: Kluwer.

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

Walley, P. (1991). *Statistical Reasoning with Imprecise Probabilities.* Chapman & Hall.

Williams, J. R. G. (2014). Decision-Making Under Indeterminacy. *Philosopher's Imprint 14*(4), 1–34.

Zynda, L. (2000). Representation Theorems and Realism About Degrees of Belief. *Philosophy of Science 67*(1), 45–69.