

HOW TO READ A REPRESENTOR

EDWARD J.R. ELLIOTT

University of Leeds

Imprecise probabilities are often modelled with *representors*, or sets of probability functions. 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 agent's 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 probability functions therein. I argue that these interpretations have some important limitations. I also propose an alternative—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.

FOR modelling rational belief, probability functions do an amazing job. Not perfect, mind you, but still there's so much they get *right*. You know what's even better than a probability function, though? A whole bunch of probability functions! Anything that can be represented by means of a single probability function can be represented with a set of such functions, plus more besides. So if we switch from modelling beliefs with probability functions over to modelling them with sets thereof—what many in philosophy call *representors*, and others call *credal sets*—then it seems we've got nothing to lose.

Well, maybe something. Aside from surface-level agreement that representors represent 'imprecise probabilities', and frequent appeals to a *credal committee* metaphor that's as apt to mislead as it is to illuminate, there really isn't much consensus on what it is that representors are supposed to represent nor how they're supposed to represent it. Worse, while everyone seems to agree that not all of the information built into a representor need reflect something psychologically *real*—a genuine property of our beliefs, as opposed to a meaningless artefact of the formalisation—competing interpretations differ regarding which properties

Contact: Edward J.R. Elliott <eelliott2@nd.edu>

of a representor have genuine representational import and which are artefacts, and hence they differ in non-trivial ways regarding what inferences can be rightly drawn about an agent's beliefs from the representation of those beliefs. A better recipe for confusion you'll not often find.

There are seven sections to this paper. After some background in §1, I will outline and discuss the two main ways of reading representors that have become especially prominent in the recent literature. In §2 and §3, I discuss *vagueness interpretations*, according to which each probability function in the representor is a 'precisification' of the subject's beliefs. In §4 and §5, I discuss *comparativist interpretations*, according to which the representor as a whole represents those comparative confidence relations that are common to all functions therein.¹ Both, I will argue, have important limitations. Finally, after an interlude on meaningfulness and measurement in §6, in §7 I present an alternative: the *functional interpretation*.

There are many potential interpretations I won't be discussing—for instance, that a representor represents higher-order uncertainty, with each function in the set corresponding to a way an agent's first-order beliefs might be given their limited introspective evidence; or that it represents the ways an ideal agent with precise beliefs might permissibly respond to inconclusive evidence. I take these to be more *epistemic* rather than *doxastic* interpretations, where the latter are my topic. But an exhaustive taxonomy of every conceivable interpretation of a set of probability functions would make for very tedious reading indeed, and would in any case be besides the point.

My goal isn't to argue that the functional interpretation is *The One True Interpretation*, nor even that it's necessarily *better* than the vagueness and comparativist interpretations. It should go without saying that there's more than one legitimate way to interpret a set of probability functions as representing *something-related-to-beliefs*, and different applications of the same formal objects in different contexts may call for different interpretations. Rather, my goal is to provide reasons for taking the functional interpretation seriously as an interesting and distinctive interpretive possibility. I've chosen to focus the earlier sections on the vagueness and comparativist interpretations in part because they are prominent, but moreover because doing so allows me to set up an illuminating contrast between important features of the functional interpretation in comparison to these better-known alternatives. Better to see what my proposal is and why it's worth considering when you can more easily compare it with what it's not.

1. I will also discuss a nearby relative of the comparativist interpretations, according to which comparative confidence relations are one among a plurality of primitive doxastic relations that are jointly represented by what's common to each of the probability functions within a representor. This *pluralist* style of interpretation is easier to introduce and explain after the comparativist interpretations have been discussed.

1. Background

Representors arose in response to concerns with the traditional single-function model of belief. Where Ω is a set of possible worlds, and propositions are subsets of Ω , we let $\mathbf{P} = \{p, q, r, \dots\}$ 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 binary unions. Then, according to the traditional model, Sally's beliefs can be represented using a single measure $\mu : \mathbf{P} \mapsto [0, 1]$ that satisfies the usual normalisation and additivity constraints—i.e., $\mu(\Omega) = 1$ and $\mu(p \cup q) = \mu(p) + \mu(q)$ whenever $p \cap q = \emptyset$.

There's something strikingly unrealistic about this. We needn't go into all of the concerns that have been raised, for they are legion—for discussion, 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 briefly 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 know neither how many are missing nor which. Now consider:

p = The global population in 2100 will be over 12 billion

q = The next card drawn from this old deck will be a heart

If you're like most people, then (a) you'll have some positive degree of confidence regarding each, and (b) you're unlikely to have *exactly* as much confidence in p as you do in q . On those assumptions, the traditional model implies that there's going to be a unique real value r such that you're exactly r times as confident in p as you are in q . So, question: exactly how much more or less confident are you in p than you are in q ?

You should find that difficult to answer, and not just because there might be some facts about the strengths of your beliefs that are introspectively hard to figure out. That may be the case, but the main problem instead seems to be that such precise values just aren't very realistic when talking about a squishy psychological quantity like strength of belief—at least not for people like us. Whatever it is about us that grounds the facts about our degrees of belief, it's not clear that there's going to be sufficient information down there to determine that we believe p down to the n^{th} degree for very large n . Indeed, it's not even obvious that p and q must stand in determinate relations of *more*, *less*, or *equal* confidence. Maybe there's no 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 *parity* is a special symmetric comparative strength relation holding only if p and q are not believed to exactly the same degree (*à la* Chang 2002).

The upshot of the example, in any case, is that there seems to be *something* about the way our beliefs are, or a way they might be, that the traditional single-function model isn't able to capture. Worry not just yet what that *something* is exactly, since that varies from author to author; worry only that it's missing when we represent beliefs with a probability function. Maybe it's several things. Either way, a more general model of belief seems to be required.

Representors to the rescue! On this new and improved approach, we represent a system of beliefs by means of a (finite or infinite, but either way non-empty) set of probability functions, $\mathcal{R} = \{\mu, \mu', \mu'', \dots\}$, all defined on the same space of propositions \mathbf{P} . When the representor contains just a single function, then it represents the very same beliefs as would have been represented by that function according to the traditional model. But when the representor contains multiple functions, on the other hand, then it represents... something else.

The *credal committee* metaphor is frequently used to give the rough idea of what that 'something else' is supposed to be. Imagine that every μ in \mathcal{R} gets a vote on what Sally's beliefs are going to be like, and the vote passes just when 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 . Likewise, if some members of the credal 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 \mathcal{R} as a whole represents neither comparative relation since the committee failed to reach unanimous agreement on the matter.

(Be warned: the metaphor is *not* an interpretation, and while it can be useful for roughly summarising how an interpretation of the model might go, it can also be misleading. There are some inferences that are naturally suggested by the metaphor that end up being licensed under some interpretations but not others. For example, unreflective application of the metaphor will suggest that Sally has more confidence in p than in q *only if* every member of her credal committee votes as such—i.e., only if $\mu(p) > \mu(q)$ for all $\mu \in \mathcal{R}$. This holds for *some* interpretations, but not all of them. Likewise, the metaphor suggests that Sally has at least as much confidence in p as she does in q only if every member of her committee votes as such—i.e., only if $\mu(p) \geq \mu(q)$ for all $\mu \in \mathcal{R}$. Again: true for some interpretations, not for all. I'll say more about this below.)

One more thing will be useful before moving on. For any representor \mathcal{R} , we define its *summary function*, \mathcal{R}^s , like so:

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

That is, $\mathcal{R}^s(p)$ picks out the set of values that the individual measures μ in \mathcal{R} assign to the proposition p . In some cases, $\mathcal{R}^s(p)$ may be an interval; in others,

$\mathcal{R}^s(p)$ may be ‘gappy’. I’ll mostly talk about cases where $\mathcal{R}^s(p)$ is an interval, but nothing hangs on this. More important to note that while summary functions are useful for 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. Distinct representors can sometimes determine the very same summary function, so in some cases there’s loss of information when going from the former to the latter. One is a set of real-valued functions, the other is a set-of-reals-valued function, and they shouldn’t be confused.

2. Vagueness Interpretations: The General Idea

Suppose one of us points towards Bruce the cat and says ‘look at Bruce!’ Presumably, there’s some indeterminacy as to what ‘Bruce’ picks out. In the vicinity of the space where we’re pointing there will be many precise cat-like things, *Bruce*₁, *Bruce*₂, *Bruce*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 we’re not referring to none of them either. Rather, one might imagine that each serves as a potential referent for ‘Bruce’ and it’s simply undecided which one it should be. Or at least that’s a plausible way to look at things. So say that each of *Bruce*₁, *Bruce*₂, *Bruce*₃, . . . , is a *precisification* of what we might mean by ‘Bruce’, the kind of thing we would be referring to if we were to somehow make our language perfectly precise. Say also, at least to begin, that anything true relative to all such precisifications of our language is true *simpliciter*, 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 very much like it). On the simplest versions of the interpretation, Sally’s representor \mathcal{R} contains all and only the probability measures μ such that, if we were to suitably precisify our language, then μ would characterise Sally’s beliefs as per the traditional model. For instance, if $\mathcal{R}^s(p) = [0.4, 0.5]$, then Sally’s degree of belief regarding p is determinately between 0.4 and 0.5 inclusive, but for any more precise degree within that interval it’s going to be indeterminate whether *that* is the degree to which Sally believes p . It would be normal in this case for fans of the vagueness interpretation to say that Sally’s beliefs are “vague over the $[0.4, 0.5]$ interval”. The presumptions, note, are that (i) each of the μ in \mathcal{R} has independent representational import, and (ii) none them determinately misrepresents Sally’s beliefs—in the sense that if $\mathcal{R} = \{\mu_1, \mu_2, \dots\}$, then according to \mathcal{R} it’s indeterminate whether μ_1 represents her doxastic state, or μ_2 does, and so on. Every function in \mathcal{R} represents its own thing, and none of them get things determinately wrong. (This will be important.)

Of course there are other ways one could flesh out the details here. 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 is true on all precisifications then it's 100% true, 0% true if it's false on all precisifications, and some middling degree of truth otherwise. There can also be variation regarding whether the vagueness is taken to result from semantic indecision, or is instead a feature of the belief system itself and independent of how we talk about it. So there isn't really one vagueness interpretation, but a family of them. The differences shouldn't matter for my purposes.

One can find examples of the vagueness interpretation in (van Fraassen 1990; 2006), (Hájek 2003), (Rinard 2015), and (Levinstein 2019). Hájek and Smithson (2012: §3) and Joyce (2010) also present what *could* be interpreted as instances of the vagueness interpretation, at least under some precisifications. One could also characterise another, strictly broader class of interpretations—the *supervaluational interpretations*—characterised by their shared application of some supervaluationist logic or other. For example, Williams (2014) interprets sets of probability functions as representing the range of precise attitudes one may rationally take towards a metaphysically indeterminate proposition. Similarly, one might take a representor to represent the rationally permissible precise belief states relative to an agent's evidence, where the facts about rational permissibility are themselves indeterminate. But these are only superficially similar to what I'm calling the vagueness interpretations. The difference is that vagueness interpretations are concerned with representing vagueness or indeterminacy relating directly to degrees of belief *themselves*, as opposed to representing vagueness or indeterminacy in connection to which precise degrees of belief *are rational*.

3. Vagueness Interpretations: The Problem

Let me start with a mention of the problem I'm *not* going to talk much about. In particular, you might worry that the simple vagueness interpretation is a bit too simple, and applying the supervaluationist rule too liberally will have some absurd consequences. After all, every member of the "credal committee" says that there's a unique real value r such that Sally is exactly r times more confident in p than she is in q , provided she has some positive degree of confidence in both. But this was precisely the sort of thing we were trying to avoid saying!

It's natural to think that such results are mere artefacts of the formalisation, an inevitable consequence of using 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' (Lewis 1999: 173). A common response, therefore, is to restrict the application of the rule when reading a representor (e.g., Zynda 2000: 49; Rinard 2017: 267). We might say that \mathcal{R} represents as true anything that's true according to every μ in \mathcal{R} , with the exception of those existential claims where no instances hold according to every such μ .

You may or may not be convinced by that response—see (Smith forthcoming) for discussion. Either way, since *that* problem closely relates to a more general and long-standing issue for supervenience that has been thoroughly discussed elsewhere, I want to pursue something different. My concern relates to those representors containing functions that, according to the traditional model, represent belief states that are very different from one another. In short, the problem here isn't so much that the μ in \mathcal{R} are precise when the goal was to represent something imprecise; the problem, rather, it is that the μ in \mathcal{R} have little in common with what they're supposed to be representing at all.

To get an initial feel for the problem, consider again the precisifications of 'Bruce'. Each of $Bruce_1$, $Bruce_2$, $Bruce_3$, and so on, has very precise boundaries, even though one might intuitively think that *Bruce* does not have precise boundaries. So there's at least one respect in which what's true for every precisification is not plausibly true of Bruce. "Not a problem", some will say, "We're not committed to saying that Bruce has precise boundaries, because we're not committed to applying the supervenience rule to every statement whatever." Grant that the response succeeds. Nevertheless—and this is the important part—in all the ways that *really* matter, every precisification of 'Bruce' is still overall very much Bruce-like. Each one walks like Bruce, each one meows like Bruce, and not a one of them, you'll observe, looks much like a cassowary. If we were to bundle up all the properties we associate with Bruce, then each of the $Bruce_i$ would satisfy the very large majority of them. They all do a good job of playing the Bruce-role, so they all have a good claim to serve as the extension of that name—that is, they all make sense *qua* precisifications of 'Bruce'. Consequently, whatever the precisifications of 'Bruce' might be, they cannot be radically unlike one another with respect to their Bruce-y properties.

Keeping that in mind, contrast two representors: $\mathcal{R}_{\text{narrow}}$ and $\mathcal{R}_{\text{wide}}$. The first, $\mathcal{R}_{\text{narrow}}$, determines only a narrow spread of values for any of the propositions in \mathcal{R} —let's say, $\mathcal{R}_{\text{narrow}}^s(p) = [0.339, 0.341]$. Now, I think we have a pretty good working idea of what Sally would be like if she were to believe such-and-such propositions to this or that precise degree. Decision theory, for example, gives us a good sense of how Sally's degrees of belief impact on her choices. Epistemology gives us a good sense of how Sally's evidence affects changes in her beliefs and hence her decisions conditional on such evidence. We have, in other words, a reasonable grasp of the main functional role associated with the systems of belief represented by probability functions as per the traditional model. And where

two such functions assign similar numerical values, they also tend to play overall similar roles. There's not a great deal of difference in *most* decision-theoretic or epistemic contexts between believing p to degree 0.339 and believing it to degree 0.341. As such, it's entirely plausible that Sally could be in a state such that her behaviour and behavioural dispositions conditional on evidence are similar to but not quite how we'd expect if she believed p to degree 0.339, and *also* similar to but not quite how we'd expect if she believed p to degree 0.341, and likewise for the many values between. The functions in $\mathcal{R}_{\text{narrow}}$ are all alike to one another, and so I can imagine that they all might represent precisified versions of a state that's simultaneously similar to all of them. Given this, the vagueness interpretation is clear enough (pun intended) for the case of $\mathcal{R}_{\text{narrow}}$.

Not so for $\mathcal{R}_{\text{wide}}$, which this time determines a much wider spread of values for many of the propositions in \mathcal{R} —say, $\mathcal{R}_{\text{wide}}^s(p) = [0, 1]$. This is often described as having beliefs that are “vague over the entire unit interval”—but what could *that* mean? As above, I have a good working idea of what Sally's behaviour (and behaviour conditional on evidence) would be like if she were absolutely certain that p . And I have a good working idea of what Sally would be like if she were absolutely certain that $\neg p$. There isn't much similarity between them. And neither is very similar to the case where Sally has 50% confidence towards p , or 25% confidence. It's hard to imagine how Sally could be in a state such that she behaves similar to but not quite how we'd expect if she believed p to degree 0, and *also* similar to but not quite how we'd expect if she believed p to degree 1, and likewise for the many values between. It seems, rather, that to be in any state such that $\mu(p) = 1$ provides a reasonable precisification thereof is *ipso facto* to be in a state such that $\mu(p) = 0$, or $\mu(p) = 0.5$, or $\mu(p) = 0.25$, doesn't.

Don't say that where Sally's beliefs are represented by $\mathcal{R}_{\text{wide}}$, then she'll be in a state that causes her to be *indeterminately disposed* between behaving in the $\mu(p) = 1$ way, the $\mu(p) = 0$ way, the $\mu(p) = 0.5$ way, and so on. For what could that really mean other than that Sally isn't really 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 —in that scenario she wouldn't be willing to pay any price for the bet! And she'll be determinately unlike what we'd expect if she had 100% confidence in p . Or 50% confidence. Or 25%.

Clarification one: the problem *isn't* that we have no account of the functional role associated with a representor like $\mathcal{R}_{\text{wide}}$. The decision-theoretic role of $\mathcal{R}_{\text{wide}}$, or its epistemic role, will be implicit in those theories in which it figures. The problem is that, *no matter what the role ends up being*, it's hard to make sense of how all the μ in $\mathcal{R}_{\text{wide}}$ can each serve as sensible precisifications of whatever it

is that $\mathcal{R}_{\text{wide}}$ supposedly represents. Those functions are associated with states that are very different from one another in all the ways that matter vis-à-vis beliefs. So at least some of the μ in $\mathcal{R}_{\text{wide}}$ will seem to determinately misrepresent Sally's beliefs, given that each has divergent implications regarding the functional role of that state that cannot all be close to the truth. Better, surely, to say that $\mathcal{R}_{\text{wide}}$ represents something *determinately* distinct from anything that might be represented by its members separately.

Clarification two: there are of course many ways we might conceivably make sense of a radically indeterminate doxastic state. Functionalists will sometimes say that an agent could be in a state that occupies the functional role of pain *for her* even while that 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 might imagine saying something like this about *believing p to degree 0* and *believing p to degree 1*. Or if you buy into quantum indeterminacy, then we could perhaps construct a *Schrödinger's believer* scenario where Sally is in a superposition of radically different belief states. No doubt there are other imaginary cases involving broken teletransporters and omnipotent demons and whatnot. 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 doxastic indeterminacy as a normal response to incomplete or non-specific evidence, not indeterminacy as a result of this one weird quirk of functionalism and hypothetical quantum mechanics experiments.

Ramsey pointed out long ago that excessive precision in the measurement of belief feels a lot like 'working out to seven places of decimals a result only valid to two' (Ramsey 1931: 76). Representors like $\mathcal{R}_{\text{narrow}}$ can capture this thought nicely. There isn't much difference, functionally, between believing p to degree 0.339 or to degree 0.341, and any plausible epistemology or decision theory is going to treat those states of belief as being generally similar to one another in most respects—and likewise the interval $[0.339, 0.341]$. So it makes sense to say that Sally's beliefs are "vague over the $[0.339, 0.341]$ interval", in much the same way it makes easy sense to speak of the boundary for 'tall' being vague over the interval from about 5'11" to 6'1". But talk of beliefs that are vague over the *entire* unit interval sounds a lot like saying the fuzziness of 'tall' extends from the tiniest infants right up to the tallest basketballers. In the case of $\mathcal{R}_{\text{wide}}$, it's a lot harder to make sense of indeterminacy as to whether this or that member of the representor represents Sally's beliefs, given that what's represented by those precisifications are all very dissimilar from one another—and, consequently, at least some of those precisifications must also represent something determinately unlike whatever they're supposedly precisifications of.

4. Comparativist Interpretations: The General Idea

According to vagueness interpretations, it's indeterminate which of the μ in \mathcal{R} are supposed to represent Sally's beliefs. According to comparativist interpretations, by contrast, in most cases when there's more than one μ in \mathcal{R} , then each of them will determinately misrepresent Sally's beliefs in some way or another. That's not a problem, because according to this kind of interpretation it's the entire set \mathcal{R} which does the representing, and no individual μ within \mathcal{R} has any representational import independent of the whole.

But I'm getting ahead of myself. I should start with *comparativism*, the idea that numerical degrees of belief are really just a way of representing what are ultimately nothing more than relations of relative confidence.² To discuss this, we'll need some more 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
 $p \nabla q$ iff Sally's confidence in p is incomparable to her confidence in q

I assume that $p \nabla q$ holds whenever p and q are not related by \succsim , \succ , or \sim in either direction, provided of course they both belong to the relevant algebra \mathbf{P} . I am therefore ignoring the possibility that there may be other non-conventional forms of comparability, such as *parity*. I 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 seems to be common ground among comparativists. (The other direction is not so obvious.) Given this, $p \not\prec q$ implies $p \not\sim q$ and $p \not\prec q$, and so it suffices from now on to say:

$p \nabla q$ iff $p \not\prec q$ and $q \not\prec p$

Given that, according to the *traditional comparativist interpretation* of a probability function μ , that function represents the facts about Sally's beliefs by virtue of representing her comparative confidences—specifically:

2. For discussion on comparativism, see (Keynes 1921), (de Finetti 1931), (Koopman 1940b; 1940a), (Fine 1973), (Zynda 2000), (Stefánsson 2017; 2018), and (Elliott 2022a; 2022b). For a recent and detailed overview on comparativism plus several connected 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 also be discussed a little later on. I've avoided this terminology because I think it's misleading: there's nothing *uniquely* measurement-theoretic about comparativism nor any nearby positions. While it is true that comparativism is traditionally founded on and explicated via the theory of fundamental extensive measurement, there are forms of measurement other than fundamental extensive measurement, and as such there are thoroughly non-comparativist perspectives which are just as 'measurement-theoretic' as any version of comparativism might have claim to be. See §6 for more on these points.

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

The important thing to note here is that there's no possibility of incomparability. The relation \geq over the real numbers is *complete* in the sense that for any two reals x, y , either $x \geq y$ or $y \geq x$; hence, any function $\mu : \mathbf{P} \mapsto [0, 1]$ interpreted as above automatically represents \succsim as being likewise complete over \mathbf{P} .

Representors provide an alternative means of representing comparative confidence relations, with the benefit of allowing for incompleteness and hence for representing incomparability. Or rather: representors provide several distinct ways of representing potentially incomplete confidence relations, corresponding to several varieties of comparativist interpretation. Again, there's not really a single interpretation here, but a family of them. One way to capture the differences is in terms of which of \succsim , \succ and/or \sim are treated as definitional primitives. On the traditional single-function model, it's typical to let \succsim be the uniquely primitive confidence relation, and simply define \succ and \sim as its asymmetric and symmetric parts respectively. This is what I did above, but it's not the only way I could have done it. I could just as easily have let \sim and \succ be the primitive relations, and then defined \succsim as the disjunction of the two (i.e., $p \succsim q$ iff $p \succ q$ or $p \sim q$). Or I could have treated \succsim and \succ as the primitives and used them to define \sim (e.g., $p \sim q$ iff $p \succsim q$ and $q \succsim p$, or iff $p \not\suc q$ and $q \not\suc p$). Or I could have let all three be considered independently primitive. The point is that *it doesn't matter*—it'll make no difference at all when it comes to reading any real-valued 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 comparativist interpretation of a representor treats \succsim as the unique primitive. On this interpretation we say that \mathcal{R} represents that $p \succsim q$ just in case every function in \mathcal{R} agrees that p is at least as probable than q , and then we let \sim and \succ be defined as the symmetric and asymmetric parts of \succsim as usual. Call this the *\succsim -interpretation*:

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

Consequence: $p \succ q$ just in case $\mu(p) \geq \mu(q)$ for all μ in \mathcal{R} , with $\mu(p) > \mu(q)$ for at least some *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, to the extent that the intended interpretation is ever explicitly and unambiguously characterised. See, for example, (Nehring 2009), (Alon & Lehrer

2014), (Miranda & Destercke 2015) (Harrison-Trainor & Holliday 2016), (Harrison-Trainor *et al.* 2018), (Konek 2019), (Ding *et al.* 2021), and (Eva & Stern 2023). We can also find a version of the \succsim -interpretation in Kaplan's 'Modest Probabilism' (1996; 2002; 2010).³

By contrast, Eva (2019: 394-5) puts forward a distinct (though obviously similar) comparativist interpretation, according to which \succ and \sim are definitionally primitive and \succsim is their disjunction. Call this the \succ/\sim -interpretation:

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

But wait—there's more! Builes *et al.* (2022) seem to put forward what we can call the \succsim/\succ -interpretation:⁴

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

It likely 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 magician, and has been asked to toss it twice. She knows that magicians' coins are often biased, though not always, and if it is biased then it'll be highly variable in which direction and to what extent. As far as she knows, it could be completely biased towards heads, or completely biased towards tails, or anything between.

3. Kaplan's several slightly different statements of 'Modest Probabilism' all presuppose an interpretation of a representor \mathcal{R} according to which (i) $p \sim q$ iff $\forall \mu \in \mathcal{R} : \mu(p) = \mu(q)$, (ii) $p \succ q$ iff $\forall \mu \in \mathcal{R} : \mu(p) \geq \mu(q)$ and $\exists \mu \in \mathcal{R} : \mu(p) > \mu(q)$, and (iii) (in Kaplan's words) you are undecided as to the relative credibility of p and q just in case $p \not\prec q$, $p \not\sim q$, and $q \not\prec p$. Assuming we can substitute ' $p \nabla q$ ' for 'you are undecided as to the relative credibility of p and q ,' and assuming as above that $p \nabla q$ implies $p \not\prec q$ and $q \not\prec p$, then (i)–(iii) are just an alternative way of formulating the \succsim -interpretation.

4. In more detail: Builes *et al.* advocate what they call the 'Comparative View', according to which $\mu \in \mathcal{R}$ iff (i) if $p \succsim q$ then $\mu(p) \geq \mu(q)$, and (ii) if $p \succ q$ then $\mu(p) > \mu(q)$. This implies that $p \succsim q$ only if $\forall \mu \in \mathcal{R} : \mu(p) \geq \mu(q)$, and likewise $p \succ q$ only if $\forall \mu \in \mathcal{R} : \mu(p) > \mu(q)$. However, it doesn't yet guarantee the converses of those two conditionals, and so we don't yet have the \succsim/\succ -interpretation. For example, suppose that $p \succsim q$ and $q \succsim r$, but $p \nabla r$. Then the Comparative View implies that for all $\mu \in \mathcal{R}$, $\mu(p) \geq \mu(q)$ and $\mu(q) \geq \mu(r)$, so also $\mu(p) \geq \mu(r)$; hence, without further assumptions, the Comparative View doesn't imply $\forall \mu \in \mathcal{R} : \mu(p) \geq \mu(r)$ only if $p \succsim r$. But if we assume that Sally's comparative confidences are rational in the sense that (a) they are or otherwise can be extended in a way that's representable by some representor, and (b) they do not have any 'gaps' that could be filled by a priori reasoning alone (e.g., if $p \succsim q$ and $q \succ r$, then it should *not* be the case that $p \nabla r$ since we should have enough to determine that $p \succ r$), then the Comparative View will imply the stronger \succsim/\succ -interpretation. It's noteworthy in any case that the Comparative View diverges from the much more common \succsim -interpretation, under which $p \succ q$ can be true even if it's not the case that $\forall \mu \in \mathcal{R} : \mu(p) > \mu(q)$.

Given this, we might decide to represent Sally’s beliefs by means of a representor, call it $\mathcal{R}_{\text{coin}}$, such that if

p = The coin will land heads on the next toss

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

then for all μ in $\mathcal{R}_{\text{coin}}$,

$$\mu(p) = \sqrt{\mu(q)}$$

and

$$\mathcal{R}_{\text{coin}}^s(p) = \mathcal{R}_{\text{coin}}^s(q) = [0, 1]$$

Don’t worry about whether you think this is the *right* way to represent Sally’s beliefs in this situation; the important point for the example is that $\mu(p) = \mu(q)$ only where $\mu(p) = 1$ or $\mu(p) = 0$, and otherwise $\mu(p) > \mu(q)$. Now, since every μ in $\mathcal{R}_{\text{coin}}$ agrees on $\mu(p) \geq \mu(q)$, but they don’t all agree on $\mu(q) \geq \mu(p)$, then according to the \succsim -interpretation we should read $\mathcal{R}_{\text{coin}}$ as saying:

$$p \succsim q, \quad p \succ q, \quad p \not\prec q, \quad p \not\bowtie q$$

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

$$p \not\prec q, \quad p \not\sim q, \quad p \not\prec q, \quad p \nabla q$$

And on the *other* other hand, the \succsim/\succ -interpretation reads $\mathcal{R}_{\text{coin}}$ as saying:

$$p \succsim q, \quad p \not\sim q, \quad p \not\prec q, \quad p \not\bowtie q$$

The foregoing is useful for highlighting the dangers arising from unreflective reliance on the credal committee metaphor. According to the \succ/\sim -interpretation and the \succsim/\succ -interpretation, every voter on the committee needs to agree that $p \succ q$ in order for $p \succ q$ to be true, in line with what the metaphor suggests, but not so for the \succsim -interpretation. Likewise, if every voter agrees that $p \succsim q$, then $p \succsim q$ according to the \succsim -interpretation and the \succsim/\succ -interpretation, but not always according to the \succ/\sim -interpretation. And while all three comparativist interpretations agree that \mathcal{R} represents $p \sim q$ just in case everyone on the committee votes $p \sim q$, they also all imply that $p \not\prec q$ inasmuch as a single voter puts their hand up for either $p \succ q$ or $q \succ p$ —which contrasts with the vagueness interpretations, according to which $p \not\prec q$ only when *every* committee member votes for $p \succ q$ or $q \succ p$. Everyone can agree that the metaphor gets some things right and some things wrong, but good luck getting them to agree on what.

5. Comparativist Interpretations: The Problem

Distinguish *comparativist interpretations* from *comparativism*. The former are just a way of reading representors, and it's uncontroversial that representors can be used to represent incomplete confidence orderings. Comparativists advocate something stronger: that's all a representor *needs* to represent, because those states of comparative confidence are what ultimately comprise our systems of belief. What matters is just the ordering: *ordinal-equivalence* is *meaning-equivalence*. That, as they say, is what's *real*; aught else is just an artefact of the numerical representation. The question for us is whether we might reasonably want representors to represent something more.

One of the most frequently cited motivations for comparativism is the plausible idea that there seems to be nothing about our beliefs that calls for a *unique* numerical representation, or any *numerical* representation at all for that matter. Comparativism is in a position to explain this, and that is an explanatory virtue of the view (see, e.g., Koopman 1940a: 269; Fine 1973: 15; Zynda 2000: 64ff; Stefánsson 2017). Builes *et al.* summarise this idea 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: 7)

A similar motivation is that comparativism is able to capture and explain certain intuitive possibilities that don't play nicely with traditional real-valued representations of belief—for instance, (a) that Sally might have more confidence in p than she does in q , without there being any particular degree to which she has more confidence in p ; and (b) that Sally's confidence in p need not be more, less, or equal to her confidence in q , since p and q may be incomparable. Fine briefly mentions something along these lines as a reason for adopting purely relational models of belief:

(2) [Comparative probability] provides a wider class of models of random phenomena than does the usual [i.e., precise, real-valued] 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. (Fine 1973: 15–6)

Now all this *would* provide a very compelling motivation for comparativism indeed, *if* comparativism were in a unique position to capture these intuitively plausible thoughts. But it isn't. Non-comparativists can say these sorts of things too, and they often do! The presupposition seems to be that the only theoretical option besides comparativism is the view that degrees of belief correspond to unique real numbers literally inscribed somewhere inside the head. In reality, though, non-comparativists just agree that the particular numbers we use are just one way among many for *numerically* representing a *qualitative* psychological system, presumably by virtue of some structural similarity. Likewise, non-comparativists just agree that the qualitative psychological systems being represented need not always have the kind of structure that allows for representation using precise real numbers. All this is common ground. What *isn't* common ground is whether the qualitative psychological systems being represented by a (precise or imprecise) numerical model of belief can be fully characterised in terms of comparative confidence relations. *That's the debate.*⁵

But there's an argument for comparativism in the nearby vicinity that's worth considering in more detail. Let *absolute degrees of belief* be the kinds of doxastic attitudes that relate an agent to a proposition and a degree (which may or may not be represented numerically); for instance, those attributed when we say that Sally is *very confident* that *p*, is *certain* that *p*, or believes *p* to *degree 0.69*, and so on. Premise one: the very notion that there are *degrees of belief* presupposes a minimal relational structure—there must at least be a transitive and reflexive ordering over them, as anything less and we'd be stretching the concept of *degrees* beyond recognition. Premise two: any system of absolute degrees of belief immediately determines a corresponding system of comparative confidences—for instance, if Sally's degree of belief in *p* is greater than her degree of belief in *q*, then $p \succ q$. Premise three: a non-transitive system of comparative confidences seems to be possible. So there seem to be possible systems of comparative confidence that correspond to no possible system of absolute degrees of belief. So absolute degrees of belief are not plausibly *more fundamental* than comparative confidence; facts about the latter do not supervene on facts about the former.

5. Here, 'qualitative' contrasts with 'numerical'. Following Tarski (e.g., 1954), let a *relational system* be understood as a set with one or more relations defined thereon. This way of using the terms is usual in the literature on measurement. The idea is that some systems—call them *numerical*—are characterised by explicit reference to numbers and numerical relations. Other systems—call them *qualitative*—can be characterised without reference to any specific numbers or explicitly numerical relations. Inasmuch as a qualitative and a numerical system share a similar relational structure, we can represent ('measure') the former by systematically mapping it into the latter. I'll say more about this in the next section. Some will want to say that 'numerical' systems are characterised wholly by their structure; hence any 'qualitative' system with the same structure instantiates that system and should also count as 'numerical' (e.g., Mitchell 2021). That might be right; but whether qualitative systems *instantiate* numerical systems or are merely *represented by* them, either way the distinction proves useful.

I'm inclined to accept the conclusion. But that's not enough yet to conclude that comparative confidences are more fundamental than absolute degrees of belief. For one thing, it may be that neither is more fundamental than the other—perhaps the facts about both fall out simultaneously from the facts about some third species of doxastic state, such as outright beliefs. Or maybe, as Lewis often suggested (e.g., 1986: 36–7; 1994: 430), we can see the system of beliefs *as a whole* as comprising the fundamental doxastic unit. On this picture—which I'll advocate in §7—comparative confidences and absolute degrees of belief both are just ways of describing salient aspects of a total doxastic state characterised by its functional role in connection to evidence and behaviour. Or, instead of positing a third state from which both comparative confidence and absolute degrees of belief derive, you might suppose that comparative confidence is merely one among a plurality of primitive kinds of doxastic state. The other primitives may include, e.g., judgements as to when propositions are evidentially independent of one another, or states of certainty, or full belief, and so on. Konek (2019: 308ff) refers to this as the *pluralist* view—that there's more than one species of primitive doxastic state, such that the facts about our beliefs cannot all be reduced just to facts about comparative confidence. Many have recommended adding at least a primitive *qualitative independence* relation alongside comparative confidence as one of the basic doxastic relations that are normally represented by our precise probabilistic models of belief (see, e.g., Domotor 1970; Fine 1973; Kaplan & Fine 1977; Luce 1978; Luce & Narens 1978), and Joyce (2010) suggests an interpretation of representors like this. (I'll say more about pluralism in the next section.)

So the explanatory motivations that are usually put forward for accepting comparativism are not very compelling—there *are* some intuitive possibilities that comparativism is able to explain, that's true, but it's not *uniquely* positioned to explain those possibilities. Moreover, a wide range of contemporary theories that employ (precise or imprecise) numerical representations of belief make regular appeal to extra-ordinal properties of those representations that *cannot* be taken to represent anything expressible wholly in terms of comparative confidence. Generally speaking, the role played by our numerical representations of belief in the bulk of current theory requires those representations to carry more than merely ordinal information. Given this, we have at least *some* reason for wanting more from an interpretation than the comparativist interpretations give us.

I cannot discuss every example in detail, but I can look at one very simple case in decision theory. Let $\mathbf{P} = \{\Omega, p, \neg p, \emptyset\}$, and suppose

$$\Omega \succ p \succ \neg p \succ \emptyset$$

According to the interpretations I've discussed, a representor \mathcal{R} will determine these comparative confidences if (but not in all cases only if), for all μ in \mathcal{R} ,

$$1 > \mu(p) > \mu(\neg p) > 0$$

This will include any \mathcal{R} such that $\mathcal{R}^s(p) \subseteq (0.5, 1)$. So according to comparativism, these representors all represent the same system of beliefs.⁶ But now take any decision theory for imprecise probabilities that generalises expected utility theory, in the sense that it includes that theory as a special case whenever \mathcal{R} is singleton. This covers all the well-known theories— Γ -maximin, E-admissibility, maximality or interval dominance. (See Troffaes 2007 for an overview.) That theory will entail that there is a decision-theoretically relevant difference between at least some, if not all, of these ‘ordinally equivalent’ representors. For instance, imagine that Sally is choosing between two gambles:

α : receive \$1 if p is true, nothing otherwise

β : receive \$2 if p is false, nothing otherwise

Which should she choose? Case 1: if $\min[\mathcal{R}^s(p)] > 2/3$, then Sally should prefer gamble α . Case 2: if $\max[\mathcal{R}^s(p)] < 2/3$, then Sally should prefer gamble β . Case 3: if $\max[\mathcal{R}^s(p)] > 2/3 > \min[\mathcal{R}^s(p)]$, then depending on the theory she might either be indifferent between the two gambles, prefer α to β , prefer β to α , or lack a preference—either way it’ll differ from either Case 1 or Case 2.

There’s nothing controversial about this. It’s well-known that decision theories for ‘precise’ probabilities usually attribute differential import to ordinally equivalent representations of belief. Preferences licensed by some probability function when conjoined with a utility function need not be licensed by another probability function when conjoined with that same utility function, even if the two probability functions determine the same confidence ordering. This is true for normative theories (e.g., expected utility theory, or risk-weighted utility theory) and descriptive theories (e.g., cumulative prospect theory). The same extends to decision theories for ‘imprecise’ probabilities. That is: pairs of representors that determine the same confidence orderings can and often do carry differential import for some decision situations, according to these theories.

And it’s no different outside of decision theory. Epistemology supplies more examples—to say nothing yet of game theory, information theory, or linguistics. Probabilistic independence is centrally important for our theories of evidence and learning, but it’s long been known that independence cannot generally be defined in terms of binary comparative confidence. (See Domotor 1970; Kaplan & Fine 1977; Luce & Narens 1978; Joyce 2010: 285ff discusses independence relations specifically in connection to representors, alongside several other epistemically important relations that cannot be formulated using comparative probability.) Theories of peer disagreement require interpersonal comparisons of confidence,

6. There will only be one *maximally inclusive* representor that corresponds to any ordering if any representor does, and comparativists may want to say that we should always use the maximally inclusive representor. The Comparative View (mentioned in fn. 4) builds in this requirement, for example. This won’t affect the point I’m making, only how it’s made.

which seem especially tricky for comparativists to explain. (See Elliott 2022b for the feasibility of interpersonal confidence comparisons within a comparativist framework.) The Principal Principle presupposes extra-ordinal distinctions between rational belief states, given that there are meaningfully distinct but ordinally equivalent objective chance functions. And so on, and on, and on. (And on.) In short: across a wide range of contexts, in decision theory and elsewhere, numerical representations of belief are attributed theoretical roles that require them to carry meaningful extra-ordinal information.

“So what?”, the inevitable interjection goes, before the argument is yet complete. “Aren’t you just presupposing that all these theories are *correct* in appealing to this extra-ordinal information, and therefore begging the question against comparativism? And don’t we have reason already to suppose these theories often help themselves to more information than they’re entitled, as for instance when standard decision theories represent decision-makers as having complete awareness of their state-space? Our current theories are rife with idealisation—so what’s to stop comparativists from simply saying that inasmuch as these theories make use of extra-ordinal information, then this is yet another idealisation not to be taken too seriously?”⁷

In response, we can focus on the forest or on the trees. Start with the trees. While it’s clearly true that current theories of belief and decision-making are over-idealising in many respects *when applied to ordinary agents*—as when they presume full awareness, or very precise gradations in degree of belief—those idealisations nevertheless seem to have a sensible interpretation in scenarios involving idealised but still *possible* agents. While *we* don’t have full awareness of our state-space, it’s not impossible that someone could. While *we* might not have very precise degrees of belief, it’s not impossible that someone could. The ‘extra information’ isn’t *meaningless*—it’s not a mere artefact of the representation, but something with a legitimate role to play in conceivable scenarios. The situation with comparativism is quite different. Comparativists aren’t saying that there’s information encoded in a probability function, or in a representor, which makes sense for ideal agents but not for us. Rather, comparativists are saying that if a theory treats two numerical models of belief differently even though they determine the same confidence orderings, then the theory is appealing to meaningless information that doesn’t have a proper interpretation for *any* scenario. And while I’ll happily concede that contemporary theories of decision-making and belief update and whatnot are unrealistic for agents like us, that fact doesn’t yet give me a reason to suppose that extra-ordinal information has no meaningful role to play.

Now the forest: the argument begs no questions against comparativism, since it neither concludes with nor is premised on anything implying the falsity of that view. Here is a summary of the whole thing, which should help:

7. I’m paraphrasing more than one commentator’s objection here.

Across a wide range of theoretical contexts, ordinarily equivalent representations of belief are typically attributed differential import. It is (at best) unclear whether our current theories can be revised so to fit with comparativism without significant loss in terms of, *inter alia*, fit with empirical data, fit with reflective intuitions about rationality, explanatory power and capacity for integration with adjacent theories—the usual theoretical virtues.⁸ So, while it's consistent with these facts that comparativism might still be correct, given the present state of theorising it is *ceteris paribus* reasonable to doubt that numerical representations of belief represent only comparative confidence orderings.

I said this earlier, and now I'll say it again: my goal is *not* to convince you that comparativism is mistaken. It's not clear it *is* mistaken. However, it's not clear that comparativism is correct either. That's the point. The goal is to provide reasons in favour of an interpretation that's capable of saying more than comparativism allows. The general schema of functional interpretation that I propose below is consistent with the possibility that comparativism is correct, but it's also consistent with the opposite possibility—and *that's a good thing*.

The fact that our current theories overwhelmingly tend to presuppose the meaningfulness of extra-ordinal information is a reason to think that such information is meaningful. It's not a conclusive reason by any means, but it's certainly not negligible either. By way of analogy, consider John Wheeler's (1964, 1980) pregeometry programme. According to Wheeler, our theories of space and time should be reconstructed in such a manner as 'breaks loose at the start from all mention of geometry and distance' (1980: 3–4)—that is, without presupposing the meaningfulness of any essentially geometric structures and concepts (even to the point of giving up the concept of *distance*), and weakening common assumptions about the nature of spacetime (such as continuity). There have been some limited attempts in this direction, such as replacing continuous spacetime in special relativity theory with a weaker discrete spacetime, with some partial successes. (See Meschini *et al.* 2004 for a user-friendly overview.) However, we're far from having anything approaching a general theory of spacetime that doesn't presuppose the meaningfulness of a very good deal of classical geometric concepts, and at this stage it's not at all clear whether such a thing is really feasible. Those facts seem to suffice for taking very seriously the possibility that certain basic geometric concepts and assumptions really *are* essential to our physical theories, and for adopting such assumptions as reasonable working hypotheses.

8. There has been some limited work towards revising decision theory such that it only employs comparative confidence relations. Fine (1973: 37ff) shows that *some* interesting decision situations can be formulated with only comparative relations. As he notes, though, 'clearly much remains to be done' (1973: 16) before we have anything that can be considered a fully adequate comparativist decision theory. That is as true today as it was then.

Back to comparativism. I've argued that, as things stand, there don't seem to be any interesting possibilities that comparativism is in a *unique* position to explain, and moreover comparativism sits ill-at-ease with the fact that numerically distinct but ordinarily equivalent representations of belief are typically afforded differential import across a very wide range of contemporary theoretical contexts. These facts may change. Maybe better theories will be developed, which do not appeal to extra-ordinal information in the numerical representation of belief. Until such time, I'm going to take the widely presumed meaningfulness of extra-ordinal information implicit in the bulk of present theorising at face value, as telling us something important about the nature of belief.

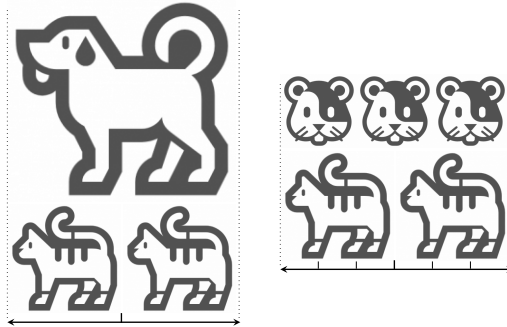
6. Beyond Doxastic Structure

Numerical models of belief are representations of some qualitative psychological system, presumably by virtue of their possessing some similarity of structure. On this I agree with the comparativists. The interesting debate is not between those who do and do not think that the numbers we happen to employ when ascribing degrees of belief are representations of some underlying psychological structure; rather, it concerns what that structure *is*. What, in other words, are the qualitative psychological 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 alternatives with an equal claim to representational adequacy?

There's two main ways we might think about these questions. On the one hand, we might think that what's being represented can be fully characterised in terms of qualitative doxastic concepts and relations. Let's refer to that as a *qualitative doxastic structure*. When comparativists say that probability functions and representors represent comparative confidence orderings, they're referring to a qualitative doxastic structure in this sense. When pluralists say that our numerical representations of belief represent some other primitive doxastic states as well, such as qualitative independence relations, they're positing a richer qualitative structure but still an essentially *doxastic* structure. A rather different approach is to suppose that the numbers represent not so much the internal structure of the belief system itself, or not *only* that, but also something about how our beliefs relate to certain other psychological phenomena—preferences, actions, evidence, for example. What's *real* may in part be a matter of the role the system of beliefs plays in our broader psychological economy.

That's all very abstract and not a little vague, and to explain it fully I'll need to take a detour through some measurement theory. Consider first the familiar story of length. Lengths are standardly measured on a ratio scale, which is to say that transformations between all the normal measures of length (meters, feet,

miles, parsecs, etc.) always preserve ratios. This is not idle stipulation: ratios of lengths on these measures have genuine physical meaning, and that meaning can be appreciated simply by considering how lengths relate to one another—that is, without considering how lengths relate to other quantities. If Spot the dog is *twice as long* as Bruce the cat, then if we were to have two copies of Bruce and line them up them head-to-tail, their combined length would be as long as Spot. And if Harry the hamster is *two-thirds as long* as Bruce, then three copies of Harry should be as long as two copies of Bruce.



A bit more formally, let $\langle \mathbf{O}, \succsim, \circ \rangle$ be the *qualitative length structure*, where \mathbf{O} is the set of physical objects, \succsim is the *at least as long* relation, and \circ is a concatenation operation with $a \circ b = c$ meaning that if a and b are lined up end-to-end then the result will be as long as c . The standard measures of length all correspond to structure-preserving mappings from the qualitative length system $\langle \mathbf{O}, \succsim, \circ \rangle$ into the numerical system $\langle \mathbb{R}^{\geq 0}, \geq, + \rangle$, where $\mathbb{R}^{\geq 0}$ is the non-negative reals and \geq and $+$ have their usual interpretations. That is, $\varphi : \mathbf{O} \mapsto \mathbb{R}^{\geq 0}$ is a structure-preserving mapping from $\langle \mathbf{O}, \succsim, \circ \rangle$ into $\langle \mathbb{R}^{\geq 0}, \geq, + \rangle$ when, for all a, b, c in \mathbf{O} ,

$$\begin{aligned}
 a \succsim b & \quad \text{iff} \quad \varphi(a) \geq \varphi(b) \\
 a \circ b = c & \quad \text{iff} \quad \varphi(a) + \varphi(b) = \varphi(c)
 \end{aligned}$$

Call this an *additive* representation, since it maps the concatenation operation into addition. Ratios are meaningful relative to additive representations of length, and that meaning is reflected directly in what’s invariant across all such representations: if φ maps $\langle \mathbf{O}, \succsim, \circ \rangle$ into $\langle \mathbb{R}^{\geq 0}, \geq, + \rangle$, then so too does ψ just in case ψ and φ are related by a ratio-preserving transformation.

The important thing to note about the example is that the underlying qualitative structure is characterised in terms of relations between lengths, without reference to any other quantities. As a result, and to put the point roughly, it is possible to explain the qualitative meaning of length ratios wholly in terms of how lengths relate to other lengths. So, if Spot is twice as long as Bruce, then that’s because $\text{Bruce} \circ \text{Bruce} = \text{Spot}$; and if Harry is two-thirds as long as Bruce, then that’s because $\text{Bruce} \circ \text{Bruce} = \text{Harry} \circ \text{Harry} \circ \text{Harry}$.

Comparativists and pluralists alike suppose that the numerical representation of belief is much like the measurement of length in this respect. On the simplest versions of traditional comparativism, the idea is that a real-valued measure μ maps $\langle \mathbf{P}, \succsim \rangle$ into $\langle \mathbb{R}^{\geq 0}, \geq \rangle$, such that $p \succsim q$ iff $\mu(p) \geq \mu(q)$. More sophisticated versions of the view add that the union of disjoint propositions behaves a lot like the concatenation operation in the case of length, and should therefore be mapped into addition: if $p \cap q = \emptyset$, then $\mu(p \cup q) = \mu(p) + \mu(q)$. Pluralists enrich the underlying qualitative system still further with additional *doxastic* relations to be captured by the numerical representation—e.g., if \perp is a qualitative independence relation, then $p \perp q$ should imply $\mu(p \cap q) = \mu(p) \cdot \mu(q)$. The ‘imprecise’ versions of these views replace the single measure with a set of measures, and then say that the qualitative relations represented are those determined in common by all measures in the set. In all these cases, what’s supposedly being represented is a *qualitative doxastic structure*, characterised in doxastic terms and without reference to other non-doxastic parts of our psychology.

But things don’t *have* to work this way. The case of length is only one model by which the measurement of belief might be understood, and it is not always possible to appreciate what a numerical model of some phenomenon represents without understanding how that phenomenon systematically interacts with others as part of a broader system. The theory of conjoint measurement was developed to explain how relations between quantities can give rise to meaningful information that’s not apparent when each is considered in isolation (Debreu 1960; Luce & Tukey 1964; Krantz *et al.* 1971). Imagine two quantities **A** and **B**, lacking in any of the apparent intrinsic structure had by the qualitative system of lengths. Still we might consider how **A** and **B** trade-off to produce varying levels in some third quantity, **C**, and from there extract meaning. For $i, j = 1, 2, \dots$, let a_i, a_j be distinct levels of **A** and b_i, b_j distinct levels of **B**. Furthermore, let \succsim now be a partial order over **C**, and let $a_i b_j$ be the level of **C** produced by a_i and b_j . Assume that **A** and **B** combine to determine **C** in an intuitively ‘additive’ way. This assumption can be rigorously characterised in purely qualitative terms, but basically amounts to a sequence of independence conditions on the structure of \succsim —for example, if $a_2 b_i \succsim a_1 b_i$ for some b_i , then $a_2 b_i \succsim a_1 b_i$ for all b_i , so the contribution to **C** made by **A** is independent of the contribution made by **B**. Given that \succsim has the appropriate structure, we can then extract an ordering over **A**: say that a_2 is more than a_1 just in case $a_2 b_i \succ a_1 b_i$ for some b_i . Moreover, we can define ratios of differences in **A**. Suppose that $a_1 b_2 \sim a_2 b_1 \succ a_1 b_1$. We read this as saying that the change from a_1 to a_2 (holding **B** fixed) produces the same effect in **C** as the change from b_1 to b_2 (holding **A** fixed); hence, the difference between $a_2 b_2$ and $a_1 b_1$ is *twice* that between $a_1 b_2$ and $a_1 b_1$, or between $a_2 b_1$ and $a_1 b_1$. Now if $a_2 b_2 \sim a_3 b_1$, then the difference between a_1 and a_3 , in terms of the contribution to **C**, is *twice* the difference between a_1 and a_2 .

These ratios in relation to \mathbf{A} have real meaning, but unlike the case of length that meaning need not correspond to any natural qualitative operations comprehensible in \mathbf{A} -terms alone—the meaning is manifest, rather, in the relationships between \mathbf{A} , \mathbf{B} and \mathbf{C} . For a conjoint measurement of a system like this, the goal will be to construct two separate functions, $\varphi_a : \mathbf{A} \mapsto \mathbb{R}$ and $\varphi_b : \mathbf{B} \mapsto \mathbb{R}$ which combine via a specified operation $f : \mathbb{R} \times \mathbb{R} \mapsto \mathbb{R}$ to determine a function $\varphi_c : \mathbf{C} \mapsto \mathbb{R}$ such that:

$$a_i b_j \succsim a_k b_l \quad \text{iff} \quad f[\varphi_a(a_i), \varphi_b(b_j)] \geq f[\varphi_a(a_k), \varphi_b(b_l)]$$

Any alternative measures in the same representational system must preserve this relation (as captured by the rule f) between the three quantities; this constrains what counts as a permissible transformation of each function individually. Think of the numerical representations of \mathbf{A} , \mathbf{B} and \mathbf{C} as a package deal; or, better, as parts of a single representation comprising three functions and a rule linking them together. What's meaningful in φ_a , then, will be tied up in how that function relates to the rest of the conjoint representation.⁹

To help illustrate this in the case of belief, consider next a well-known example from Lyle Zynda (2000). Let actions be represented in the usual way as a functions (α, β, \dots) from states ($s_i, i = 1, 2, \dots$) to consequences, and let \succsim_p be Sally's preference relation. We assume that Sally's preferences are a function of her beliefs about the state of the world and the desires she has in relation to the consequences of her actions. A *decision-theoretic representation* of this system will be a conjoint representation consisting in a representation of beliefs, a representation of desires, and a decision rule by which they jointly determine a system of preferences. An *expected-utility representation*, for example, will comprise a probability function μ and a utility function v such that

$$\alpha \succsim_p \beta \quad \text{iff} \quad \sum \mu(s_i) v[\alpha(s_i)] \geq \sum \mu(s_i) v[\beta(s_i)]$$

Now, if the probability-utility pair (μ, v) represents \succsim_p in this manner, then so too does (μ, v^*) , where

$$v^*(c) = 9v(c) + 1$$

Since v and v^* are related by a linear (interval-preserving) transformation, and since utilities are typically understood to be measured on an interval scale (like temperatures in degrees Celsius) the usual response to this fact is that there's no meaningful difference between the two functions. They represent the same

9. See (Krantz *et al.* 1971: 17–20) for more details on the present example. There are, of course, many other conjoint structures than the one I've (very briefly!) outlined here. See also (Kahneman & Tversky 1979) for an early application of the theory of additive conjoint measurement in decision theory, whereby they establish that both utilities and decision weights (roughly: beliefs plus risk attitudes) can be measured on ratio scales under the assumption that they pairwise determine preferences as described by prospect theory.

information in different ways—what matters is what’s invariant between them. And since v and v^* do *not* have ratios in common, we should say that ratios are *not* meaningful in the measurement of desire. So far so good. What Zynda notes is that whenever an expected-utility representation (μ, v) exists, then so too will another decision-theoretic representation (μ^*, v) , where

$$\mu^*(p) = 9\mu(p) + 1$$

The catch is that this time we need to adjust the decision rule by which μ^* and v jointly determine \succsim_p :

$$\alpha \succsim_p \beta \text{ iff } \sum \mu(s_i)v[\alpha(s_i)] - [\alpha(s_i)] \geq \sum \mu(s_i)v[\beta(s_i)] - [\beta(s_i)]$$

Call this a *valuation-maximising representation*. If any expected-utility representation (μ, v) of \succsim_p exists, then a valuation-maximising representation (μ^*, v) of \succsim_p also exists, and vice versa. And by analogy with v and v^* , one might be tempted to infer from this something about *meaningfulness* in μ and μ^* —namely, that there’s no meaningful difference between the two functions, that what matters is what’s invariant. Consequence: ratios are not meaningful in μ , since they’re not invariant across μ and μ^* . As Zynda suggests,

One might point out that μ^* is simply a linear transformation of μ , and argue that in the case of probabilities (like utilities and temperatures) this is a difference that makes no difference. This approach commits... to taking as real properties of degrees of belief at most those properties that are common to *both* [μ and μ^*].... 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 probabilities, being purely qualitative... The concept of degree of belief on this strategy becomes a *purely ordinal notion*... (2000: 64–5, notation altered for consistency)

But there were some leaps there. While the example *does* highlight something important about meaningfulness in μ , *this* is very much not it.

First note that while μ and μ^* share their ordinal structure, that’s not *all* they share. The linear transformation linking μ and μ^* preserves lots of properties, not just the ordering. Most importantly, the transformation is bijective, so $\mu(p) \neq \mu(q)$ iff $\mu^*(p) \neq \mu^*(q)$ and consequently if $\mu_1 \neq \mu_2$ then $\mu_1^* \neq \mu_2^*$. And in just the same way that differences between ordinally equivalent but non-identical probability functions μ_1 and μ_2 can make a difference for our preferences according to the expected utility rule, differences between ordinally equivalent but non-identical μ_1^* and μ_2^* will likewise matter according to the valuation-maximisation rule. The same will necessarily be true for any possible ‘redefinition’ of μ . So the example cannot support treating the concept of *degree of belief* as ‘a purely ordinal notion’ after all—extra-ordinal information still matters.

The reader may note that $\mu(p) \neq \mu(q)$ iff $\mu^*(p) \neq \mu^*(q)$ precisely because linear transformations preserve ratios of differences. *But do not place any weight on this fact*, for therein lies the deeper error. Let

$$\mu^\dagger(p) = \mu^*(p)^2 = 81\mu(p) + 18\mu(p) + 1$$

The transformation from μ to μ^\dagger , or from μ^* to μ^\dagger , does not preserve difference ratios. Nevertheless, whenever Sally's preferences \succsim_p can be given an expected-utility representation (μ, v) , or a valuation-maximising representation (μ^*, v) , then they can also be given a *schmaluation-maximising representation* (μ^\dagger, v) such that

$$\alpha \succsim_p \beta \quad \text{iff} \quad \sum [\sqrt{\mu^\dagger(s_i)} - 1] v[\alpha(s_i)] \geq \sum [\sqrt{\mu^\dagger(s_i)} - 1] v[\beta(s_i)]$$

In fact, we can even construct 'equivalent' decision-theoretic representations where not even *orderings* are preserved. For *any* transformation that takes us from μ to some other μ^* , then provided the transformation is bijective and therefore invertible, there will be at least one (potentially very complicated) rule by which they can be combined to generate the same preferences relative to (μ^*, v) as the expected utility rule generates relative to (μ, v) . So there's approximately *nothing* that's preserved across *all* the belief functions that might figure in one or another decision-theoretic representation—aside from the utterly trivial requirement that different degrees of belief must be assigned different values.

The lesson here is that *meaningfulness* in the representation of any quantity is only sensibly defined relative to a fixed choice of representational format.¹⁰ Ratios of lengths are meaningful *when lengths are represented additively*—that is, when the qualitative length system $\langle \mathbf{O}, \succsim, \circ \rangle$ is represented in $\langle \mathbb{R}^{\geq 0}, \geq, + \rangle$ —but additive representations are only one among infinitely many ways in which we might choose to measure length. Hölder (1901) showed that the system of lengths can be given a multiplicative representation in the system $\langle \mathbb{R}^{\geq 1}, \geq, \times \rangle$, and ratios are never invariant across multiplicative representations. Nor will

10. This can be made more precise as follows. Where \mathbf{X} and \mathbf{Y} are any sets, R_1, R_2, \dots are relations defined on \mathbf{X} , and S_1, S_2, \dots are relations defined on \mathbf{Y} , suppose that there exists at least one structure-preserving mapping from the relational system $\mathcal{X} = \langle \mathbf{X}, R_1, R_2, \dots \rangle$ into the system $\mathcal{Y} = \langle \mathbf{Y}, S_1, S_2, \dots \rangle$. Further, where S is any n -ary relation on \mathbf{Y} , let $R(S, \varphi)$ be the relation induced on \mathbf{X} by S under φ , in the sense that $(x_1, \dots, x_n) \in R(S, \varphi)$ iff $(\varphi(x_1), \dots, \varphi(x_n)) \in S$. Then we say that S is \mathcal{X} -*meaningful relative to* \mathcal{Y} exactly when $R(S, \varphi)$ doesn't depend on the particular choice of mapping: $R(S, \varphi) = R(S, \psi)$ for any other structure-preserving mapping ψ from \mathcal{X} into \mathcal{Y} . This amounts to saying that S is meaningful relative to a choice of representational format (i.e., choice of representational system \mathcal{Y}) whenever it corresponds to the *same* relation on \mathbf{X} regardless of how we choose to represent the system \mathcal{X} within \mathcal{Y} . So ratios are meaningful in additive measures of length, since the meaning of those ratios doesn't depend on the particular choice of additive scale. For more discussion on meaningfulness, see (Pfanzagl 1968), (Luce 1978), (Narens 1985), and especially (Luce *et al.* 1990).

additive and multiplicative representations have any ratios in common. In fact, approximately *nothing* is preserved across *all* possible numerical representations of length, aside from the utterly trivial requirement that different lengths must be assigned different values. Not even orderings are preserved across all measures of length in that sense, but it would be absurd to conclude that orderings are therefore meaningless.

Likewise for decision-theoretic representations. The fact that ratios vary between μ and μ^* implies nothing whatsoever about the meaningfulness of those ratios, because expected utility representations and valuation-maximising representations involve distinct representational formats. Stronger: what Zynda-style examples actually establish is that ratios in μ really are meaningful *relative to expected-utility representations*, precisely because any transformation of μ that does *not* preserve ratios must therefore employ a different form of decision rule. But the real trick here is in recognising that the qualitative meaning of those ratios need not be expressible in purely doxastic terms. When we're modelling beliefs in a decision-theoretic context, the psychological structure we're trying to represent is not necessarily something internal to system of beliefs itself, considered in isolation from anything else and characterised in purely doxastic terms, but instead at least partly something about the relations that hold between beliefs, desires, and preferences. That is why we cannot alter 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 in the production of preferences. Ratios really *are* meaningful in the measurement of belief—at least according to expected utility theory—but we should not presume their meaning can be fully captured in purely qualitative doxastic terms and without reference to the role beliefs play as part of a broader psychological system.

This lesson has long been appreciated in the case of utilities. From Ramsey (1931) through von Neumann & Morgenstern (1944) to Savage (1954), the orthodox account of why difference ratios in utility functions are meaningful has appealed to the role that desirabilities play as part of a broader system. Considered wholly in isolation, there's no immediate reason to suppose that our desires should be measured on anything stronger than an ordinal scale: one desires *this* more than *that*. It's when those desires interact with beliefs in the production of preferences under conditions of uncertainty that the need for a richer measure is manifest. Two utility functions may be ordinally equivalent, but if they vary in their difference ratios then they'll be differentiated in at least some decision situations—and therein lies the qualitative meaning of those difference ratios. Given the intimate connection between desires and beliefs, it's a mystery that we should have been inclined to treat the representation of the beliefs any differently.

7. Functional Interpretations

With so much setup, I can in this final section be mercifully brief. The concern with the vagueness interpretations was that a representor \mathcal{R} sometimes seems to represent a doxastic state that's determinately unlike what's represented by any of the μ in \mathcal{R} . Comparativist interpretations agree on this point: where \mathcal{R} represents an incomplete confidence ordering, then every μ in \mathcal{R} will determinately misrepresent that ordering. But comparativist interpretations do not play nicely with contemporary theories, which overwhelmingly tend to presuppose the meaningfulness of extra-ordinal information. Pluralists do strictly better on that front, since they allow that \mathcal{R} may carry additional representational import beyond the comparative confidence orderings it determines. However, pluralists still presuppose that the psychological structures underwriting our numerical representations of belief—the structures that ultimately explain what is and is not meaningful in those representations—are non-conjoint, purely doxastic qualitative structures. And that's not obvious either.

So here's my thought: if it may end up being impossible to appreciate what's *real* versus what's *artifact* in a formal model of belief without appreciating the role those models play in the psychological theories that make use of them, then why not just take those roles themselves to be what's real? We do not *have* to come up with an interpretation of representors that's independent of the theories in which they figure, since the interpretation of \mathcal{R} relative to a psychological theory—of decision making, say, or a theory of belief update, or better still a theory that combines both—can just be *the thing that plays the \mathcal{R} -role in that theory*.

In more detail, suppose that T is some decision-theory-cum-epistemology in which representors have a role to play. In the usual functionalist manner (à la Lewis 1970), we treat representors as theoretical terms implicitly defined by their role within T . According to T itself, the state that's designated by a representor \mathcal{R} always perfectly occupies the \mathcal{R} -role that T sets out. But T might be mistaken, such that nothing perfectly occupies the \mathcal{R} -role even if something still comes close to doing so. Thus we take the *meaning* of \mathcal{R} relative to T to be a function from worlds to whatever it is that does the best job of satisfying the \mathcal{R} -role at that world, if anything does, and provided it does so well enough. The *extension* of \mathcal{R} relative to T will be whatever the meaning designates at our world. Two theories T and T' will typically determine distinct meanings for \mathcal{R} , and in that sense different interpretations of \mathcal{R} ; though they may also be associated with the same interpretation in the sense of fixing on the same extension for \mathcal{R} .

Two representors \mathcal{R} and \mathcal{R}' are meaningfully distinct *according to T* just in case \mathcal{R} and \mathcal{R}' play distinct roles within that theory. On a (hypothetical) comparativist decision theory, \mathcal{R} and \mathcal{R}' ought to play the same role just in case they determine the same confidence relations. For the theories we actually have, this won't be

true. Of course, whether we actually ought to treat \mathcal{R} and \mathcal{R}' as designating distinct doxastic states depends on what we take the most plausible theories to be—there's not much point worrying about whether \mathcal{R} and \mathcal{R}' are meaningfully distinct according to this or that theory, if we don't have much reason to think those theories are plausible. Thus, I claim, we have reason to treat \mathcal{R} and \mathcal{R}' as meaningfully distinct *simpliciter* inasmuch as our *best* theories of rational belief and decision-making posit distinctive roles for \mathcal{R} and \mathcal{R}' .

Eriksson & Hájek (2007: 204ff) once proposed something much like what I have in mind here. What they propose is that 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 and intuitions about rational belief and decision-making. They called their view *primitivism*, but they also note that their proposal is very much in the spirit of functionalism—the main difference being that the functionalist will want to say that our theories *implicitly define* what 'degrees of belief' are via their distinctive roles, whereas they question whether this should really be counted as a 'definition' (see Eriksson & Hájek 2007: 210). 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 mine.

But not quite the same. Eriksson & Hájek's proposed primitives are *absolute degrees of belief*. That makes sense inasmuch as we're modelling beliefs in the traditional manner, since everything a probability function says about a total belief state can be derived from what it says about the particular degree of belief it associates with each proposition. But when dealing with representors, we'd be wise not to take absolute 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}^s does, but a summary function can omit information relevant to the role played by the representor it summarises. \mathcal{R}_{coin} assigns the very same maximally imprecise interval to p and to q , but it would be a mistake to say that Sally's attitudes towards p and q are the same. Better instead to let the entire system of beliefs be our primitive, represented by \mathcal{R} , and characterise that total system of beliefs by the functional role played by \mathcal{R} in the best theories we have that make use of such models.

The important point is that the functional interpretation carries no presupposition that meaningful differences between the systems of belief represented by \mathcal{R} and \mathcal{R}' must be explicable by reference to purely doxastic qualitative structures—in terms of comparative confidences, say, give or take some other doxastic relations, and without reference to the relations between belief and the

rest of our psychology. We can still say, of course, that if $\forall \mu \in \mathcal{R} : \mu(p) > \mu(q)$ and $\forall \mu \in \mathcal{R}' : \mu(p) < \mu(q)$, then \mathcal{R} represents greater confidence in p than q while \mathcal{R}' represents the reverse. The functional interpretation is not committed to the meaninglessness of such relations—quite the opposite. But it's not committed to saying that *everything* meaningful in a representor can be expressed in a similar fashion. Sometimes, the most we might be able to say in purely qualitative terms is that \mathcal{R} and \mathcal{R}' just represent *different* systems of belief—as evidenced by their distinctive roles in the production of preferences for instance, or how they give rise to divergent choice behaviour conditional on evidence. That is what sets the functional interpretation apart, and it's a thing worth having.

8. Acknowledgements

I started on this paper in 2016, though it was quite different back then. Thanks are due to Miriam Bowen, Seamus Bradley, Jakob Donskov, Jessica Isserow, and Robbie Williams discussion on various bits and pieces of one version or another. Thanks also to audiences at the University of Leeds, the University of Notre Dame, and the CUNY Cognitive Science Speaker Series. Finally, thanks to two referees for *Ergo*, for their constructive and useful feedback.

References

- Alon, Shiri and Lehrer, Ehud (2014). Subjective multi-prior probability: A representation of a partial likelihood relation. *Journal of Economic Theory*, 151:476–492.
- Bradley, Seamus (2014). Imprecise Probabilities. *Stanford Encyclopedia of Philosophy*.
- Builes, David, Horowitz, Sophi, and Schoenfield, Miriam (2022). Dilating and contracting arbitrarily. *Nous*, 56(1):3–20.
- Chang, Ruth (2002). The Possibility of Parity. *Ethics*, 112(4):659–688.
- de Finetti, Bruno (1931). Sul Significato Soggettivo Della Probabilità. *Fundamenta Mathematicae*, 17(1):298–329.
- Debreu, Gerard (1960). Topological methods in cardinal utility theory. In Arrow, K. J., Karlin, S., and Suppes, P., editors, *Mathematical Methods in the Social Sciences*, pages 16–26. Stanford University Press.
- Ding, Yifeng, Holliday, Wesley H., and Icard III, Thomas F. (2021). Logics of Imprecise Comparative Probability. *International Journal of Approximate Reasoning*, 132:154–180.
- Domotor, Zoltan (1970). Qualitative information and entropy structures. In Hintikka, J. and Suppes, P., editors, *Information and inference*, pages 148–194. Reidel.
- Elliott, Edward (2022a). Comparativism and the Measurement of Partial Belief. *Erkenntnis*, 87:2843–2870.
- Elliott, Edward (2022b). What is 'Real' in Interpersonal Comparisons of Confidence. *Australasian Journal of Philosophy*, 100(1):102–116.
- Eriksson, Lina and Hájek, Alan (2007). What are degrees of belief? *Studia Logica*, 86(2):183–213.

- Eva, Benjamin (2019). Principles of Indifference. *The Journal of Philosophy*, 116(7):390–411.
- Eva, Benjamin and Reuben, Stern (2023). Comparative Opinion Loss. *Philosophy and Phenomenological Research*, 107(3):613–637.
- Fine, Terrence L. (1973). *Theories of Probability: An Examination of Foundations*. Academic Press.
- Fishburn, Peter C. (1986). The axioms of subjective probability. *Statistical Science*, 1(3):335–345.
- Hájek, Alan (2003). What Conditional Probability Could Not Be. *Synthese*, 137(3):273–323.
- Hájek, Alan and Smithson, Michael (2012). Rationality and indeterminate probabilities. *Synthese*, 187:33–48.
- Harrison-Trainor, Matthew, Holliday, Wesley H., and Icard III, Thomas F. (2018). Inferring Probability Comparisons. *Mathematical Social Sciences*, 91:62–70.
- Harrison-Trainor, Matthew, Holliday, Wesley H., and Icard, Thomas F. (2016). A Note on Cancellation Axioms for Comparative Probability. *Theory and Decision*, 80(1).
- Hölder, Otto (1901). Die Axiome der Quantitat und die Lehre vom Mass. *Ver. Verh. Kgl. Saschisis. Ges. Wiss. Leipzig, Math.-Phyisc. Classe*, 53:1–63.
- Jeffrey, Richard (1983). Bayesianism with a human face. *Testing Scientific Theories, Minnesota Studies in the Philosophy of Science*, 10:133–56.
- Joyce, James (2005). How Probabilities Reflect Evidence. *Philosophical Perspectives*, 19:153–178.
- Joyce, James (2010). A Defense of Imprecise Credences in Inference and Decision Making. *Philosophical Perspectives*, 24(1):281–323.
- Kahneman, Daniel and Tversky, Amos (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–291.
- Kaplan, Mark (1996). *Decision Theory as Philosophy*. Cambridge University Press.
- Kaplan, Mark (2002). Decision theory and epistemology. In Moser, P., editor, *The Oxford Handbook of Epistemology*, pages 434–462. Oxford University Press, Oxford.
- Kaplan, Mark (2010). In defense of modest probabilism. *Synthese*, 176(1):41–55.
- Kaplan, Mark and Fine, Terrance L. (1977). Joint orders in comparative probability. *The Annals of Probability*, 5:161–179.
- Keynes, John M. (1921). *A Treatise on Probability*. Macmillan, New York.
- Konek, Jason (2019). Comparative Probabilities. In *The Open Handbook of Formal Epistemology*, pages 267–348. PhilPapers Foundation.
- Koopman, Bernard O. (1940a). The Axioms and Algebra of Intuitive Probability. *Annals of Mathematics*, 41(2):269–292.
- Koopman, Bernard O. (1940b). The Bases of Probability. *Bulletin of the American Mathematical Society*, 46(10):763–774.
- Krantz, David H., Luce, R. Duncan, Suppes, Patrick, and Tversky, Amos (1971). *Foundations of measurement, Vol. I: Additive and polynomial representations*. Academic Press.
- Levinstein, Ben (2019). Imprecise Epistemic Values and Imprecise Credences. *Australasian Journal of Philosophy*, 97(4):741–760.
- Lewis, David (1970). How to Define Theoretical Terms. *The Journal of Philosophy*, 67(13):427–446.
- Lewis, David (1986). *On the Plurality of Worlds*. Cambridge University Press.
- Lewis, David (1994). Reduction of Mind. In Guttenplan, S., editor, *Companion to the Philosophy of Mind*, pages 412–431. Blackwell.
- Lewis, David (1999). Many, but almost one. In *Papers in Metaphysics and Epistemology*, pages 164–182. Cambridge University Press.

- Luce, R. Duncan (1978). Dimensionally Invariant Numerical Laws Correspond to Meaningful Qualitative Relations. *Philosophy of Science*, 45(1):1–16.
- Luce, R. Duncan, Krantz, David H., Suppes, Patrick, and Tversky, Amos (1990). *Foundations of Measurement, Vol. III: Representation, Axiomatization, and Invariance*. Dover, New York.
- Luce, R. Duncan and Narens, Louis (1978). Qualitative independence in probability theory. *Theory and Decision*, 9:225–239.
- Luce, R. Duncan and Tukey, John W. (1964). Simultaneous conjoint measurement: a new scale type of fundamental measurement. *Journal of Mathematical Psychology*, 1(1):1–27.
- Meschini, Diego, Lehto, Markku, and Piilonen, Johanna (2004). Geometry, pregeometry and beyond. *Stud.Hist.Philos.Mod.Phys.*, 36:435–464.
- Michell, Joel (2021). Representational Measurement Theory: Is its number up? *Theory & Psychology*, 31:3–23.
- Miranda, Enrique and Destercke, Sebastien (2015). Extreme points of the credal sets generated by comparative probabilities. *Journal of Mathematical Psychology*, 64-65:44–57.
- Narens, Louis (1985). *Abstract Measurement Theory*. MIT Press.
- Nehring, Klaus (2009). Imprecise probabilistic beliefs as a context for decision-making under ambiguity. *Journal of Economic Theory*, 144:1054–1091.
- Pfanzagl, Johann (1968). *Theory of Measurement*. Wiley, New York.
- Ramsey, Frank P. (1931). Truth and probability. In Braithwaite, R. B., editor, *The Foundations of Mathematics and Other Logical Essays*, pages 156–198. Routledge, London.
- Rinard, Susanna (2015). A Decision Theory for Imprecise Credences. *Philosopher's Imprint*, 15(7):1–16.
- Rinard, Susanna (2017). Imprecise Probability and Higher Order Vagueness. *Res Philosophica*, 94(2):257–273.
- Savage, Leonard J. (1954). *The Foundations of Statistics*. Dover.
- Seidenfeld, Teddy (1988). Decision Theory Without ‘Independence’ or Without ‘Ordering’. *Economics and Philosophy*, 4(2):267–290.
- Smith, Nicholas J. J. Interpreting Imprecise Probabilities. *Philosophical Quarterly*.
- Stefansson, H. Orri (2017). What Is ‘Real’ in Probabilism? *Australasian Journal of Philosophy*, 97(3):573–587.
- Stefansson, H. Orri (2018). On the Ratio Challenge for Comparativism. *Australasian Journal of Philosophy*, 96(2):380–390.
- Sturgeon, Scott (2008). Reason and the grain of belief. *Nous*, 42:139–165.
- Tarski, Alfred (1954). Contributions to the theory of models I. *Indagationes Mathematicae*, 16:26–32.
- Troffaes, Matthias C.M. (2007). Decision making under uncertainty using imprecise probabilities. *International Journal of Approximate Reasoning*, 45:17–29.
- van Fraassen, Bas (1990). Figures in a Probability Landscape. In Dunn, J. and Gupta, A., editors, *Truth or Consequences*, pages 345–356. Kluwer, Amsterdam.
- van Fraassen, Bas (2006). Vague Expectation Value Loss. *Philosophical Studies*, 127:483–491.
- von Neumann, John and Morgenstern, Oskar (1944). *Theory of Games and Economic Behavior*. Princeton University Press.
- Walley, Peter (1991). *Statistical Reasoning with Imprecise Probabilities*. Chapman & Hall.
- Wheeler, John A. (1964). Geometrodynamic and the issue of the final state. In De Witt, C. and

- De Witt, B. S., editors, *Relativity, groups and topology*, pages 317–520. Gordon and Breach.
- Wheeler, John A. (1980). Pregeometry: Motivations and prospects. In *Quantum theory and gravitation*. Academic Press.
- Williams, J. Robert G. (2014). Decision-Making Under Indeterminacy. *Philosopher's Imprint*, 14(4):1–34.
- Zynda, Lyle (2000). Representation Theorems and Realism About Degrees of Belief. *Philosophy of Science*, 67(1):45–69.