

# How to Read a Representer

Authors Removed

## Abstract

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 agents beliefs would be if any vagueness were precisified away; and comparativist interpretations, according to which the set represents those comparative confidence relations that are common to all functions therein. We argue that these interpretations come with important limitations. We 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.

**Keywords:** Imprecise probabilities · Representors · Supervaluationism · Comparative probability · Measurement · Functionalism

## Introduction

For modelling rational belief, probability functions are amazing. Not perfect, of course, but 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 with 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 nothing to lose.

Well, maybe something. Besides some 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 exactly that representors are supposed to represent nor how they're supposed to represent it. Worse, while everyone seems to agree that not all the information built into a representor need reflect something psychologically *real*—a genuine property of our beliefs, as opposed to a mere artefact of the formalisation—competing interpretations differ regarding which properties of a representor have genuine representational import and which are meaningless artefacts, and hence they differ in non-trivial ways regarding what inferences can be rightly drawn about an agent's beliefs from a representor model of those beliefs. A better recipe for confusion you'll not often find.

There are five sections to this paper. After some background in §1, we outline and discuss two ways of reading representors that have become especially prominent in the recent literature: *vagueness interpretations* in §2, and *comparativist interpretations* in §3. Both have important limitations. Finally, after an interlude on meaningfulness and measurement in §4, we present an alternative interpretation in §5—the *functional interpretation*.

Of course there are many interpretations we don't discuss. 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 multiple ways an ideal agent with precise beliefs might permissibly respond to inconclusive evidence. We take these to be more epistemic rather than doxastic interpretations, where the latter are our topic. But an exhaustive taxonomy of every conceivable interpretation would make for very tedious reading indeed, and such would be besides the point in any case. It should go without saying that there's more than one potential 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. We've no intention of saying otherwise.

Our goal is not to argue that the functional interpretation is *The One True Interpretation*, nor even that it's necessarily *better* than the vagueness and comparativist interpretations; rather, it's to provide reasons for taking the functional interpretation seriously as an interesting and distinctive interpretive possibility. We've focused on the vagueness and comparativist interpretations in §2 and §3 because they are particularly prominent, and moreover because doing so lets us set up an illuminating contrast between important features of the functional interpretation in comparison to more common alternatives. Better to understand what our proposal is and why it's worth considering when you can more easily compare it with what it's not.

## 1. Background

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

There's something strikingly unrealistic about this. We needn't go into all of the concerns that have been raised, for they are legion—see (Jeffrey 1983), (Seidenfeld 1988), (van Fraassen 1990), (Walley 1991), (Kaplan 1996), (Joyce 2005; 2010), (Sturgeon 2008), (Hájek 2012), (Alon and Lehrer 2014), and (Bradley 2014). But it won't hurt to consider one example (adapted from Fishburn 1986). Imagine that before you sits an old pack of cards. You've been told that some of the cards are missing, but that's all you're told—you don't know how many are missing nor which. 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 you'll have some positive degree of confidence regarding each, and you're unlikely to have *exactly* as much confidence in  $p$  as you do in  $q$ . On those assumptions, the traditional model implies there's a unique real value  $r$  such that you're exactly  $r$  times as confident in  $p$  as you are in  $q$ . So 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 may be some facts about the strengths of your beliefs that are introspectively hard to determine (though that may also be true). The problem instead seems to be that such precise values simply aren't very realistic when measuring a squishy psychological quantity like strength of belief, at least for people like us. Whatever it is about us that grounds the facts about our degrees of belief, there just isn't sufficient information down there to determine that we believe  $p$  down to the  $n^{\text{th}}$  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 determinate fact of the matter as to which one you believe more; or maybe they're determinately incomparable. You might even be sympathetic to the idea that your strength of belief in  $p$  can be *on a par* with your strength of belief in  $q$ , where *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 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 what that something is just yet, worry only that it's missing. Maybe it's several things. Either way, a more general model of belief is evidently required.

Representors to the rescue! On this new and improved approach, we should 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. When the representor contains multiple functions, then it represents... something else.

What that 'something else' should be taken to be is rarely spelled out in much detail, and where it is the specifics vary from person to person. Still, the *credal committee* metaphor is frequently employed to give the rough idea. Imagine that every  $\mu$  in  $\mathcal{R}$  gets a single vote on what Sally's beliefs are going to be like, and the vote passes just when the committee is unanimous. If every  $\mu$  votes that Sally's confidence in  $p$  is greater than her confidence in  $q$ , then Sally's confidence in  $p$  really is greater than her confidence in  $q$ —even if there's no precise value  $r$  such that all members of the committee agree that Sally's confidence in  $p$  is exactly  $r$  times her confidence in  $q$ . Similarly, if some members of the 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 since the committee failed to reach 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 many 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 credal committee metaphor will suggest that Sally has more confidence in  $p$  than in  $q$  only if every member of her committee votes as such—i.e.,  $\mu(p) > \mu(q)$  for all  $\mu \in \mathcal{R}$ . This holds for some of the interpretations we consider below, but not all. 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.,  $\mu(p) \geq \mu(q)$  for all  $\mu \in \mathcal{R}$ . Again: works for some interpretations, not for all.)

A tiny bit more notation and terminology will be useful before moving on to the interpretations. For each representor  $\mathcal{R}$ , define its *summary function* 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  $\mu$  in  $\mathcal{R}$  assign to  $p$ . In some cases,  $\mathcal{R}^s(p)$  may be an interval; in others,  $\mathcal{R}^s(p)$  may be ‘gappy’. We’ll mostly describe cases where  $\mathcal{R}^s(p)$  is an interval, but nothing hangs on this. More important to note that while summary functions can be good for, e.g., describing the spread of values assigned to a proposition by the ‘credal committee’, a summary function is *not* just another way of representing a representor. Distinct representors can determine the very same summary function, so in some cases there’s loss of information if we replace representors with their summary functions. 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

Suppose one of us points at 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 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 much like it. On the simplest versions, Sally’s representor  $\mathcal{R}$  contains all and only the probability measures  $\mu$  such that, if we were to suitably precisify

our language, then  $\mu$  would characterise Sally’s beliefs as per the traditional model. For instance, if  $\mathcal{R}^s(p) = [x, y]$ , then Sally’s degree of belief regarding  $p$  is determinately between  $x$  and  $y$  inclusive, but for any more precise degree within that interval it’s indeterminate whether *that* is the degree to which Sally believes  $p$ . It’s common in this case to for adopters of the vagueness interpretation to that Sally’s beliefs are “vague over the  $[x, y]$  interval”. The presumption, note, is that each of the  $\mu$  in  $\mathcal{R}$  has representational import, and none of them determinately misrepresents Sally’s beliefs: if  $\mathcal{R} = \{\mu_1, \mu_2, \dots\}$ , then it’s indeterminate whether  $\mu_1$  represents her doxastic state, or  $\mu_2$  does, and so on, but none of the functions in  $\mathcal{R}$  determinately get things wrong. (This will be important.)

Of course there are other ways we could flesh out the details. Traditional supervaluationism says that truth is truth-under-all-precisifications, and something that’s true on some precisifications but false on others will simply lack a truth-value. Degree-theoretic supervaluationism says that if something 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 results from some kind of 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 our 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. We could also define another, broader category of interpretation—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 we’re calling the vagueness interpretation. The difference is that the vagueness interpretation is concerned with representing vagueness or indeterminacy relating directly to degrees of belief *themselves*, as opposed to representing vagueness or indeterminacy in connection to which degrees of belief *are rational*.

#### *Limitations with vagueness interpretations*

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!

It's natural to suppose that these sorts of 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. (1999, p. 173)

A common response is therefore to restrict the rule when reading a representor (e.g., Zynda 2000, p. 49; Rinard 2017, p. 267). We might say that  $\mathcal{R}$  represents as determinately true anything that's true according to every  $\mu$  in  $\mathcal{R}$ , with the exception of those existential claims where no instances are true according to every such  $\mu$ .

You may or may not be convinced by that response—see (Smith forthcoming) for critical discussion. Either way, since *that* particular problem relates to a much more general and long-standing issue for supervaluationism that has been thoroughly discussed elsewhere, we want to pursue something different. Our concern relates specifically 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 in these cases isn't so much that the  $\mu$  in  $\mathcal{R}$  are precise when the goal was to represent something imprecise—the problem rather is that the  $\mu$  in  $\mathcal{R}$  have little in common with what they're supposed to be representing at all.

To get a feel for the problem, consider again the precisifications of 'Bruce'. Each of Bruce<sub>1</sub>, Bruce<sub>2</sub>, Bruce<sub>3</sub>, and so on, have very precise boundaries, even though one might intuitively think Bruce himself does not have such precise boundaries. So there's at least one respect in which what's true for every precisification is not plausibly true of Bruce himself. "Not a problem", some will say, "We don't have to say that Bruce has precise boundaries because we shouldn't be applying the supervaluationist rule to every statement whatever." Grant the response succeeds. Nevertheless—and this is the important part—in all the ways that *really* matter, every precisification of 'Bruce' is still very much Bruce-like. Each one walks like Bruce, each one meows like Bruce, and not a one of them, you'll observe, looks very much like a cassowary. If we were to bundle up all the properties that we associate with Bruce, then each of Bruce<sub>1</sub>, Bruce<sub>2</sub>, Bruce<sub>3</sub>, ..., would satisfy the very large majority of them. They all do a very 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  $\mathbf{P}$ —e.g.,  $\mathcal{R}_{\text{narrow}}^s(p) = [0.339, 0.341]$ . Now we have a pretty good 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 working sense of how Sally's degrees of belief impact on her choices. Epistemology gives us a

good working 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 functional role associated with the beliefs represented by the different probability functions as per the traditional model. And where two such functions assign very similar numerical values, they tend to play overall very 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, and likewise for any values inbetween.

Given this, it makes good sense to say that Sally’s degree of belief towards  $p$  is “vague over the [0.339, 0.341] interval”, since any precise value within that range might not be *quite* right but would still do a good enough job overall of explaining and predicting her behaviour. That is, it’s 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 values between. The functions in  $\mathcal{R}_{\text{narrow}}$  are all like one another, and as such we can imagine they all might represent precisified versions of a state that’s simultaneously similar to all of them. So we’re happy to accept that the vagueness interpretation is clear enough (pun intended) in the case of  $\mathcal{R}_{\text{narrow}}$ .

Now compare  $\mathcal{R}_{\text{wide}}$ , which determines a wide spread of values for many of the propositions in  $\mathbf{P}$ —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, we 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 we 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. So it’s quite 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 all the many values between. 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, for example, that where Sally’s beliefs are given 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 possibly 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  $p$ —such an agent wouldn’t be willing to pay anything for the bet. And she’ll be determinately unlike what we’d expect if she were 100% confident that  $p$ . Or 50% confident. Or 25%.

A clarification: the problem *isn't* that we cannot have any account of the 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, and distinct theories will attribute distinct roles. The problem is that, *no matter what the role ends up being*, it's hard to make sense of how all the very different  $\mu$  in  $\mathcal{R}_{\text{wide}}$  can each serve as sensible precisifications of whatever it is that  $\mathcal{R}_{\text{wide}}$  represents. Those precise functions are associated with states that are very different from one another in all that ways that matter vis-à-vis beliefs. So at least some of the  $\mu$  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 what's represented by its members.

Additional clarification: there are of course many ways we might potentially make sense of a radically indeterminate doxastic state. Functionalists will sometimes say that an agent could conceivably 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 could imagine someone 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* situation where Sally is in a superposition of radically different belief states. No doubt there are other cases involving 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 vague evidence, not indeterminacy as a result of this one weird quirk of functionalism or 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' (1931, p. 76). Representors like  $\mathcal{R}_{\text{narrow}}$  capture this thought nicely. There isn't much difference, functionally, between believing  $p$  to degree 0.339, or 0.341, and any plausible epistemology or decision theory is going to treat those states of belief as 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 some cases, at least, it's hard to make sense of vagueness as to whether this or that member of the representor represents Sally's beliefs, given that what's represented by those precisifications are very dissimilar from one another, and, consequently, at least many of those precisifications must also represent something dissimilar to whatever state they're supposedly precisifications of.



### 3. Comparativist Interpretations

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

But we're getting ahead of ourselves. We should start with *comparativism*, the idea that numerical degrees of belief are just a way of representing what are ultimately nothing more than relations of relative confidence.<sup>1</sup> To discuss this we'll need some notation:

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

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

$$p \nabla q \text{ iff } p \not\prec q \text{ and } q \not\prec p$$

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

$$\left\{ \begin{array}{ll} p \succsim q & \text{iff } \mu(p) \geq \mu(q) \\ p \succ q & \text{iff } p \succsim q \text{ and } q \not\prec p \\ p \sim q & \text{iff } p \succsim q \text{ and } q \succsim p \end{array} \right\}$$

Note that there's no possibility of incomparability here. The  $\geq$  relation on the reals is complete, in that for any two real values  $x, y$ , either  $x \geq y$  or  $y \geq x$ ; hence, any real-valued function  $\mu$  automatically represents  $\succsim$  as being likewise complete over  $\mathbf{P}$ .

---

<sup>1</sup> For discussion on comparativism and similar positions: (Keynes 1921), (de Finetti 1931), (Koopman 1940b; 1940a), (Fine 1973), (Zynda 2000), (Stefánsson 2017; 2018), and (Elliott 2022a; 2022b). For a recent overview on comparativism and related topics, see (Konek 2019). In Konek's terminology, the position being discussed at present is the 'unary measurement-theoretic view'; the 'pluralist measurement-theoretic' version will be discussed briefly later on. We note there's nothing uniquely *measurement-theoretic* about comparativism or any nearby views. Comparativism is typically founded on the theory of fundamental extensive measurement, but there's much more to measurement theory than fundamental extensive measurement and there's no shortage of non-comparativist views that are as 'measurement-theoretic' as any version of comparativism might have claim to be. See §4 for more discussion.

Representors provide us with an alternative means of representing comparative confidence relations, with the added benefit of allowing for incompleteness and hence for representing incomparability. Or more accurately: representors provide us with several distinct ways of representing potentially incomplete confidence relations, corresponding to several varieties of comparativist interpretation. Again, there’s not a single interpretation here, but a family of them.

One way to capture the difference between comparativist interpretations 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. But this isn’t the only way we could do things. We could just as easily let  $\sim$  and  $\succ$  be the primitive relations, and define  $\succsim$  as the disjunction of the two (i.e.,  $p \succsim q$  iff  $p \succ q$  or  $p \sim q$ ). Or we could treat  $\succsim$  and  $\succ$  as our primitives and use them to define  $\sim$  (e.g.,  $p \sim q$  iff  $p \succsim q$  and  $q \succsim p$ , or iff  $p \not\succ q$  and  $q \not\succ p$ ). Or we could just let all three be considered independently primitive. The point is that *it doesn’t matter*—it’ll make no difference when it comes to reading any real-valued function  $\mu$  as a representation of Sally’s comparative confidences. But these choices do make a difference when we shift over to representing incomplete relations on 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:

$$\left\{ \begin{array}{l} p \succsim q \text{ iff } \forall \mu \in \mathcal{R} : \mu(p) \geq \mu(q) \\ p \succ q \text{ iff } p \succsim q \text{ and } q \not\succ p \\ p \sim q \text{ iff } p \succsim q \text{ and } q \succsim p \end{array} \right\}$$

Consequence:  $p \succ q$  just in case  $\mu(p) \geq \mu(q)$  for all  $\mu$  in  $\mathcal{R}$ , with  $\mu(p) > \mu(q)$  for at least some *but not necessarily all* of them. This is 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 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 *forthcoming*). We can also find the  $\succsim$ -interpretation in Kaplan’s ‘Modest Probabilism’ (1996; 2002; 2010).<sup>2</sup>

<sup>2</sup> 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) you are undecided as to the relative credibility of  $p$  and  $q$  just in case  $p \not\succ q$ ,  $p \not\sim q$ , and  $q \not\succ p$ . 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  $p \not\succ q$  and  $q \not\succ p$ , then Kaplan’s (i)–(iii) are just another way of expressing the  $\succsim$ -interpretation.

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

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

But wait—there’s more! Builes *et al.* (2022) seem to put forward what we can call the  $\succsim/\succ$ -interpretation:<sup>3</sup>

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

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 a number of times. 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 it could be completely biased towards heads, or completely biased towards tails, or anything between. Suppose

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

We might decide to represent Sally’s beliefs by means of a representor  $\mathcal{R}_{\text{coin}}$  such that for all  $\mu$  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  $\mu(p) > \mu(q)$ .

---

<sup>3</sup> In more detail: Builes *et al.* advocate what they call the ‘Comparative View’, according to which  $\mu \in \mathcal{R}$  iff, if  $p \succsim q$  then  $\mu(p) \geq \mu(q)$ , and if  $p \succ q$  then  $\mu(p) > \mu(q)$ . The Comparative View implies that  $p \succsim q$  only if  $\forall \mu \in \mathcal{R} : \mu(p) \geq \mu(q)$ , and  $p \succ q$  only if  $\forall \mu \in \mathcal{R} : \mu(p) > \mu(q)$ , however without further assumptions we cannot guarantee the converses of those two conditionals. Suppose, e.g., 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  $\mu(p) \geq \mu(r)$ ; hence, without some more assumptions about  $\succsim$ , the Comparative View doesn’t imply  $\forall \mu \in \mathcal{R} : \mu(p) \geq \mu(r)$  only if  $p \succsim r$ . But assume that Sally’s comparative confidences are rational in the sense that they can be extended in a way that’s representable by some representor, and they do not have any ‘gaps’ that could be filled by a priori reasoning alone (e.g., if  $p \succsim q$  and  $q \succsim r$ , then it should *not* be the case that  $p \nabla r$  since we should have enough to determine that  $p \succsim r$ ). Then, the Comparative View will imply the stronger  $\succsim/\succ$ -interpretation. In any case, the Comparative View diverges from the more common  $\succsim$ -interpretation, which allows that  $p \succ q$  even if not  $\forall \mu \in \mathcal{R} : \mu(p) > \mu(q)$ .

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

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

On the other hand, since neither  $\mu(p) > \mu(q)$  nor  $\mu(p) = \mu(q)$  for all  $\mu$ , 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:

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

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, as the metaphor suggests, but not so for the  $\succsim$ -interpretation. Likewise, if every committee member 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 imply that  $p \not\sim q$  inasmuch as a single voter puts their hand up for either  $p \succ q$  or  $q \succ p$ —in contrast to the vagueness interpretations, which implies  $p \not\sim q$  only when  $\forall \mu \in \mathcal{R} : \mu(p) \neq \mu(q)$ . Apparently, everyone can agree that the credal committee metaphor gets some things right and some things wrong, but good luck getting them to agree on what.

#### *Limitations with comparativist interpretations*

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

This question becomes especially pressing when we recognise that contemporary theories and models that employ numerical representations of belief make regular appeal to extra-ordinal properties of those representations that cannot be taken to represent anything expressible in terms of mere comparative confidence. This includes descriptive and normative theories of the structure of belief, of evidence and learning, decision-making, and more, both for precise and imprecise degrees of belief, in philosophy and other disciplines. We cannot discuss every example in detail, for they are many and various, but we can look at one very simple decision-theoretic example. Let  $\mathbf{P} = \{\Omega, p, \neg p, \emptyset\}$ , and suppose

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

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

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

This will include any  $\mathcal{R}$  such that  $\mathcal{R}^s(p)$  is any sub-interval of  $(0.5, 1)$ . So according to comparativism, they all represent the same system of beliefs.<sup>4</sup> 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 when  $\mathcal{R}$  is singleton. This covers all the well-known theories— $\Gamma$ -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 representors. For instance, imagine Sally is choosing between two gambles:

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

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

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

There’s nothing controversial about any of this. It’s well-known that contemporary decision theories for ‘precise’ probabilities attribute differential import to ordinally-equivalent numerical representations of belief. That is: preferences licensed by some probability function when conjoined with a given utility function need not be licensed by a different probability function when conjoined with those same utilities, even if the two determine the very same confidence ordering. This is true for normative theories (such as expected utility theory, or risk-weighted utility theory) and descriptive theories (such as cumulative prospect theory). And it’s likewise straightforward to see that the same extends to contemporary decision theories for ‘imprecise’ probabilities. That is: pairs of representors that determine the same confidence orderings can and often do carry differential decision-theoretic import according to these theories. In general, the role numerical representations of belief play in contemporary theories of decision-making requires them to carry meaningful information going beyond what can be expressed in terms of comparative confidence relations.

Likewise for a great deal of work outside decision theory. Epistemology supplies more examples—to say nothing yet of game theory, information theory, philosophy of language, or linguistics. Probabilistic independence is centrally important for our theories of evidence and learning, but cannot be defined in terms of binary comparative confidence. (See Domotor 1970; Kaplan & Fine

---

<sup>4</sup> There will only be one maximally inclusive representor that represents any ordering, and comparativists may want to say that we should always use the maximally inclusive  $\mathcal{R}$ . The Comparative View (discussed in fn.3) builds in this requirement, for instance. This doesn’t affect the point we’re making, only how it’s made.

1977; Luce & Narens 1978.) Joyce (2010, pp.285ff) discusses the case of independence specifically in connection to representors, as well as several other epistemically important relations that cannot be formulated using comparative probability. Epistemic utility theory appeals to properties that differentiate ordinally-equivalent probability functions. (See Mayo-Wilson & Wheeler 2019, p.19, for an example.) Theories of peer disagreement require interpersonal comparisons of confidence, which seem particularly difficult for comparativists to capture. (See Elliott 2022b for discussion on the feasibility of interpersonal confidence comparisons within a comparativist framework.) The Principal Principle presupposes extra-ordinal distinctions between rational belief states, in light of the fact that there are meaningfully distinct but ordinally-equivalent objective chance functions. In short: across a wide variety of contexts, 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 making appeal to this additional information, and therefore begging the question against comparativism? And don’t we have some reason already to suppose these theories often help themselves to more information than they’re entitled, as for instance when most contemporary decision theories represent decision-makers as having complete awareness of their state-space, or when they employ infinitely precise distinctions in degree of belief? Our current theories are rife with idealisation—what’s to stop comparativists from simply responding that inasmuch as these theories do make use of some extra-ordinal information, then that is yet another idealisation not to be taken too seriously?”<sup>5</sup>

In response, we can focus on the forest or on the trees. Start with the trees. While it’s clear that current theories of belief and decision-making are unrealistic and over-idealising in many respects *when applied to ordinary agents*—as when they presume full awareness, for instance—those same idealisations still seem to have a perfectly sensible interpretation in conceivable scenarios involving idealised agents. While *we* might not have full awareness of our state-space, an appropriately idealised agent presumably could. The ‘extra information’ in these cases isn’t *meaningless*—it’s not a mere artefact, but something that has a legitimate role to play (albeit in limited applications). The situation with comparativism is quite different. Comparativists aren’t saying that there’s further information encoded in a probability function, or in a representor, which makes sense for idealised agents but not us. If a theory treats  $\mu$  and  $\mu'$  differently, or  $\mathcal{R}$  and  $\mathcal{R}'$  differently, despite those fixing the same confidence orderings, then that isn’t just the theory being inapplicable to mere mortals—it’s the theory making appeal to information that doesn’t have an interpretation for any scenario no matter how idealised. We’ll happily accept that contemporary theories of decision-making, say, or of belief update, are often somewhat unrealistic when applied to ordinary agents, but we don’t yet see a reason to suppose they’re fundamentally unsound in application to ideal cases.

---

<sup>5</sup> We are paraphrasing more than one reader’s comments here, so we’re going to dwell on the objection in some detail—maybe too much, but it’s a point that comes up frequently.

Now the forest: all these many theories *may* be mistaken in appealing to extra-ordinal information, but if you think we're begging the question then you've misunderstood the structure of the argument that we're going for. That argument is not yet complete, but the 'question begging' objection needs to be nipped in the bud. Here's a summary of the whole thing that should help:

Across a range of theoretical contexts, the overwhelming tendency is that ordinally-equivalent representations of belief are often attributed differential import. It is at best unclear whether these theories can be revised 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. While it's possible that comparativism is correct and these theories should all be revised as such, given the present state of things it is *ceteris paribus* reasonable to doubt that numerical representations of belief represent no more than comparative confidence orderings.

Our goal is *not* to convince you that comparativism is mistaken. It's not clear that it *is* mistaken. But it's not clear it's correct either, and that's the point. The goal is to motivate an interpretation of representors that's capable of saying more than the comparativist interpretations afford; for this we need only establish that there are considerable reasons to want more. Comparativism may end up being the correct view, and the scheme of interpretation we'll propose below is consistent with that possibility—but it's also consistent with the opposite possibility. And that's a good thing.

In any case, the argument begs no questions against comparativism, since it neither concludes with nor is premised on anything implying the falsity of that view. We're going to take it for granted that the first step of the argument has been established: across a wide variety of theoretical contexts relating to belief, it's generally presumed that ordinal-equivalence is not meaning-equivalence. Given that, should one wish to deny our conclusion there's really only three potential avenues for response. We consider each in turn.

First: one might deny the principle underlying the argument. That principle goes something like this:

Assume (i)  $p$  is a widespread presupposition that does theoretical work, (ii) it's unclear whether denying  $p$  can be achieved without significant loss, and (iii) there are no compelling reasons for accepting  $\neg p$  even if doing so would lead to significant loss. Then it is not irrational to have some doubt as to  $\neg p$ , nor irrational to accept  $p$  as a working hypothesis.

We'd be amazed if that principle, or something to the same effect, were not widely accepted. (Maybe *that's* question begging too!) Consider 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, pp. 3–4)—that is, without presupposing the meaningfulness of 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 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, at present 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. More importantly, it’s not at all clear whether such a thing really is feasible. That fact alone seems to suffice for taking very seriously the possibility that certain basic geometric assumptions really *are* essential to our physical theories, and for adopting those assumptions as reasonable working hypotheses.

Second: one might show that the relevant theories can be revised or replaced with a theory that’s consistent with comparativism, without significant loss. Suffice it to say that comparativists at this stage simply haven’t shown that such is possible, and it remains at best unclear whether it is. That’s not to say there’s been no work in this direction—but it is to say that any such work is far from complete. Fine (1973, pp. 37ff), for example, shows that *some* interesting decision situations can be formulated within a comparativist context. As he notes, though, ‘clearly much remains to be done’ (1973, p. 16) before we have anything that can be considered an adequate normative decision theory for comparativism. That is as true today as it was back then. Nor have we any particular reason to suppose that comparativist decision theories can compete with current non-comparativist theories for empirical adequacy. Likewise, and as already mentioned, comparativists have yet to put forward a plausible account interpersonal confidence comparisons, which seem *particularly* tricky to account for on their view. And so on, and on, and on.

Third: one might try to argue that there are independent reasons supporting comparativism so compelling that we ought rationally to adopt the position *even if* revising our current theories accordingly would lead to some significant loss. Are there any such arguments? It is surprisingly difficult to find *any* clear arguments for comparativism, but those that do exist tend to focus on its capacity to explain certain ideas and intuitive possibilities. One of the most frequently cited motivations for comparativism is the idea that there’s nothing about our beliefs that calls for a *unique* numerical representation, or indeed any *numerical* representation at all (see, e.g., Koopman 1940a, p. 269; Fine 1973, p. 15; Zynda 2000, pp. 64ff; Stefánsson 2017). Builes *et al.* summarise it nicely:

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



And this *would* indeed provide compelling reasons for comparativism, if that view were uniquely positioned to capture that thought. But non-comparativists can say these things too, and they often do! The implicit presupposition seems to be that the only option besides comparativism is the view that degrees of belief correspond to unique real numbers literally inscribed somewhere inside the head, when in reality non-comparativists can readily agree that the particular numbers we use are just one way to *numerically* represent what must fundamentally be a *qualitative* psychological system—presumably by virtue of structural similarity. All this is common ground.

(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 on that set. The idea is that some systems—call them *qualitative*—can be characterised without reference to any specific numbers or specifically numerical relations. Other systems—call them *numerical*—are characterised by reference to specific numbers and numerical relations. Where qualitative and numerical systems share a similar structure, we can represent the former by systematically mapping it into the latter. We’ll say more about this shortly. This way of using ‘qualitative’ and ‘numerical’ is common in the literature on measurement. 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., Michell 2021). That might be right, but the distinction proves useful whether qualitative systems *instantiate* numerical systems or are merely *represented by* them.)

A better argument appeals not to the representational nature of the numbers by which we happen to ascribe graded beliefs, but rather to the possible doxastic states that comparativism might be in a special position to explain. Fine briefly mentions something like this:

(2) [Comparative probability] provides a wider class of models of random phenomena than does the usual quantitative theory. [...] Point (2) refers to the curious phenomenon that there exist relatively simple examples of what we consider to be valid [comparative probability] statements that are incompatible with any representation in the usual quantitative theory. (Fine 1973, pp. 15–6)<sup>6</sup>

Now it’s clear there are some things comparativists are in a position to explain that cannot be represented in the usual quantitative theory. Comparativism allows for the possibility that Sally might have more confidence in  $p$  than in  $q$  without there being any particular degree to which she is more confidence, and it allows for incompleteness in the confidence ordering. *These* claims are not the special province of comparativism, though, and non-comparativists of various stripes and colours can also say that the psychological structures underlying the numerical representation of belief may not always support precise representation on the real number line.

---

<sup>6</sup> The ‘wider’ claim is not quite correct. As there are distinct probability functions (and representors) that determine identical confidence relations, the class of comparative probability models cannot be wider than the class of precise (or imprecise) numerical models—just different.

However, the following version of the argument is worth taking rather more seriously. Let *absolute degrees of belief* be the kinds of attitudes that relate an agent to a proposition and a degree that may or may not be represented numerically. These will be the attitudes attributed when we say that Sally is *very confident* that  $p$ , or believes  $p$  to *degree 0.69*, and so on. Premise one: the very notion of *degrees* of belief presupposes that those degrees have a minimal relational structure: they must at least be transitive and reflexive, as anything less and we'd be stretching the usual concept of *degrees* beyond recognition. Premise two: a system of absolute degrees of belief always determines a corresponding system of comparative confidences—e.g., if Sally's degree of belief in  $p$  is greater than her degree of belief in  $q$ , then  $p \succ q$ . Premise three: non-transitive comparative confidences seem to be possible. Conclusion: the facts about comparative confidence cannot reduce to the facts about absolute degrees of belief.

The third premise isn't obvious but it is plausible, and as such the argument provides some reason to suppose that absolute degrees of belief are not more fundamental than comparative confidences. But that's not enough yet to conclude that the facts about our beliefs supervene on the facts about comparative confidence. One might suppose, for instance, that comparative confidence is merely one species of primitive doxastic state among others. These others may include, e.g., qualitative judgements of probabilistic independence, or judgements that one proposition is evidence for another, or states of certainty, or full belief, or some other states with a belief-*ish* flavour. Konek (2019, pp. 308ff) refers to this as the *pluralist* view. Joyce (2010) seems to advocate an interpretation of representors like this, and many have recommended adding at least a qualitative independence relation alongside comparative confidence (see, e.g., Domotor 1970; Fine 1973; Kaplan & Fine 1977; Luce 1978; Luce & Narens 1978).

Indeed, the argument doesn't even provide sufficient reason to conclude that comparative confidence is more fundamental than absolute degrees of belief. It may turn out that neither is more fundamental than the other—perhaps the facts about both fall out simultaneously from the facts about some third kind of state, such as outright beliefs. Or maybe, as Lewis often suggested (e.g., 1986, pp. 36–7; 1994, p. 430), we can see the system of beliefs *as a whole* as comprising the fundamental doxastic unit. On this picture—which we'll be advocating in §5—our talk of comparative confidences and absolute degrees of belief are ultimately just ways of describing salient aspects of a single total doxastic state characterised by its functional role in relation to evidence and behaviour. For the present it doesn't matter which is correct; what matters is they're still on the table.


Let's sum up. As things presently stand, ordinally-equivalent representations of belief are typically afforded differential import across a wide range of theoretical contexts, including decision theory, epistemology, and more. It is at best uncertain whether we can treat ordinal-equivalence as meaning-equivalence without engendering some significant loss, and there appear to be no reasons that would force the rational acceptance of comparativism despite that uncertainty. These facts may change. In the meantime, we're going to take the presumed meaningfulness of extra-ordinal information in our present theories at face value, as telling us something important about the nature of belief.

#### 4. Beyond Doxastic Structure

Numerical models of belief are best thought of as representations of some fundamentally qualitative psychological system, presumably by virtue of their possessing some similarity of structure. On this we agree with the comparativists. The interesting debate is not between those who do and do not think the numbers we happen to employ when ascribing degrees of belief are representations—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 representational alternatives with an equal claim to adequacy?

There's two main ways we might think about that structure. On the one hand, what's being represented might be fully characterised in terms of qualitative doxastic relations and concepts. Let's refer to that as a *qualitative doxastic structure*. When comparativists say that probability functions and representors represent comparative confidence orderings, they're refer to a qualitative doxastic structure in this sense. When pluralists say that our numerical representations of belief also capture 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 qualitative 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, so to explain it fully we'll need to take a detour through some measurement theory. We consider first the familiar story of length. Lengths are standardly measured on a ratio scale, and thus transformations between normal conventional measures of length (feet, miles, meters, parsecs, etc.) 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—and without considering how lengths relate to any 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  $\langle \mathbf{O}, \succsim_l, \circ \rangle$  be the *qualitative length structure*, where  $\mathbf{O}$  is the set of physical objects,  $\succsim_l$  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 is as long as  $c$ . Standard measures of length correspond to structure-preserving mappings from  $\langle \mathbf{O}, \succsim_l, \circ \rangle$  into the 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,  $\varphi$  maps  $\langle \mathbf{O}, \succsim_l, \circ \rangle$  into  $\langle \mathbb{R}^{\geq 0}, \geq, + \rangle$  when  $\varphi : \mathbf{O} \mapsto \mathbb{R}^{\geq 0}$  and for all  $a, b, c \in \mathbf{O}$ ,

1.  $a \succsim_l b$  iff  $\varphi(a) \geq \varphi(b)$
2.  $a \circ b = c$  iff  $\varphi(a) + \varphi(b) = \varphi(c)$

Call this an *additive* representation. Ratios are meaningful relative to additive representations of length, and that meaning is reflected in what's invariant across all such representations: if  $\varphi$  maps  $\langle \mathbf{O}, \succsim_l, \circ \rangle$  into  $\langle \mathbb{R}^{\geq 0}, \geq, + \rangle$ , then so too does  $\psi$  iff  $\psi$  and  $\varphi$  are related by a ratio-preserving transformation.

The thing to note is that the underlying structure is characterised in terms of qualitative relations between lengths, without reference to other quantities. In that sense, we can explain the qualitative meaning of length ratios wholly in terms of how lengths relate to other lengths: if Spot is twice as long as Bruce and Harry two-thirds as long as Bruce, then that's *because*  $\text{Bruce} \circ \text{Bruce} = \text{Spot}$  and  $\text{Bruce} \circ \text{Bruce} = \text{Harry} \circ \text{Harry} \circ \text{Harry}$ . Comparativists and pluralists alike suppose that the representation of belief is essentially similar to the measurement of length in this respect. On the simplest versions of traditional comparativism, the idea is that  $\mu$  maps  $\langle \mathbf{P}, \succsim \rangle$  into  $\langle \mathbb{R}^{\geq 0}, \geq \rangle$ , so  $p \succsim q$  iff  $\mu(p) \geq \mu(q)$ . More sophisticated versions add that the union of disjoint propositions behaves a lot like a concatenation operation 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 system further with additional *doxastic* relations to be captured in the numerical representation—e.g., if  $\perp$  is a qualitative independence relation, then  $p \perp q$  should imply  $\mu(p \cap q) = \mu(p) \cdot \mu(q)$ . In all cases, the system being represented is characterised in qualitative doxastic terms, without reference to the relations between beliefs and other non-doxastic parts of our psychology. But things don't *have* to work this way. It is not always possible to appreciate what a numerical model of some phenomenon represents without understanding how that phenomenon 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  $\mathbf{A}$  and  $\mathbf{B}$  lacking in any of the apparent intrinsic structure had by the qualitative system of lengths. Still we might consider how  $\mathbf{A}$  and  $\mathbf{B}$  trade-off to produce varying levels in some third quantity,  $\mathbf{C}$ . For  $i, j = 1, 2, \dots$ , let  $a_i, a_j$  be distinct levels of  $\mathbf{A}$  and  $b_i, b_j$  distinct levels of  $\mathbf{B}$ . We let  $\succsim_c$  be a partial order over  $\mathbf{C}$ , and we let  $a_i b_j$  be the level of  $\mathbf{C}$  produced by  $a_i$  and  $b_j$ . Assume that  $\mathbf{A}$  and  $\mathbf{B}$  can each be measured on at least an interval scale, and that they combine in an intuitively 'additive' way. These assumptions can be rigorously characterised in purely qualitative terms, but basically amount to a sequence of independence conditions on the structure of  $\succsim_c$ —for example,

if  $a_2b_i \succsim_c a_1b_i$  for some  $b_i$ , then  $a_2b_i \succsim_c a_1b_i$  for all  $b_i$ , so the contribution to  $\mathbf{C}$  made by  $\mathbf{A}$  is independent of the contribution made by  $\mathbf{B}$ . Given that  $\succsim_c$  has the appropriate structure, we can then extract an ordering over  $\mathbf{A}$ : say that  $a_2$  is more than  $a_1$  just in case  $a_2b_i \succ_c a_1b_i$  for some  $b_i$ . Moreover, we can define ratios of differences in  $\mathbf{A}$ . Suppose that for some  $a_1b_1$  less than  $a_2b_1$ ,  $a_1b_2 \sim_c a_2b_1$ . We read this as saying that the change from  $a_1$  to  $a_2$  (holding  $\mathbf{B}$  fixed) produces the same effect in  $\mathbf{C}$  as the change from  $b_1$  to  $b_2$  (holding  $\mathbf{A}$  fixed); hence, the difference between  $a_2b_2$  and  $a_1b_1$  is *twice* that between  $a_1b_2$  and  $a_1b_1$ , or between  $a_2b_1$  and  $a_1b_1$ . Now if  $a_2b_2 \sim_c a_3b_1$ , then the difference between  $a_1$  and  $a_3$ , in terms of the contribution to  $\mathbf{C}$ , is *twice* the difference between  $a_1$  and  $a_2$ .

The ratios in relation to  $\mathbf{A}$  have real meaning, but unlike the case of length that meaning need not correspond to any natural qualitative relations expressible in  $\mathbf{A}$ -terms alone—they are manifest, rather, in the relationships between  $\mathbf{A}$ ,  $\mathbf{B}$  and  $\mathbf{C}$ . In a conjoint measurement of this kind of system, the goal will be to construct two separate measures,  $\varphi_a : \mathbf{A} \mapsto \mathbb{R}$  and  $\varphi_b : \mathbf{B} \mapsto \mathbb{R}$  which compose to determine a numerical measure  $\varphi_c : \mathbf{C} \mapsto \mathbb{R}$  according to a specified operation  $f : \mathbb{R} \times \mathbb{R} \mapsto \mathbb{R}$  such that

$$a_ib_j \succsim_c a_kb_l \text{ iff } f[\varphi_a(a_i), \varphi_b(b_j)] \geq f[\varphi_a(a_k), \varphi_b(b_l)]$$

Any alternative representations into the same representational system must preserve this relation between the three quantities; this constrains what counts as a permissible transformation of each function individually.<sup>7</sup> Think of the numerical representations of  $\mathbf{A}$ ,  $\mathbf{B}$  and  $\mathbf{C}$  as a package deal; or, better, as parts of a single representational system comprising three functions and a numerical operation linking them together. What's meaningful in  $\varphi_a$ , then, will be tied up in how that function relates to the rest of the conjoint representation.

To help illustrate this in the case of belief, consider a well-known example from Lyle Zynda (2000). Let actions be represented in the usual way as a functions from states ( $s_i, i = 1, 2, \dots$ ) to consequences, and let  $\succsim_p$  be Sally's preference relation; we assume her 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 an operation by which they jointly determine a system of preferences. An *expected-utility representation*, for example, will comprise a probability function  $\mu$  and a utility function  $v$  such that  $\alpha \succsim_p \beta$  iff

$$\sum \mu(s_i)v[\alpha(s_i)] \geq \sum \mu(s_i)v[\beta(s_i)]$$

Now, if  $(\mu, v)$  represents  $\succsim_p$  in this manner, then so too does  $(\mu, v^*)$ , where

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

---

<sup>7</sup> See Krantz *et al.* (1971, pp. 17–20) for more details on the present example, which involves additive conjoint measurement. There are, of course, vastly many other conjoint structures than the one we've very briefly described here. See also Kahneman & Tversky (1979) for an 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.

Since utilities are usually understood to be measured on an interval scale—like temperatures in degrees Celsius—the typical response to this fact is that there’s no meaningful difference between  $v$  and  $v^*$ . They represent the same thing in different ways; what matters is what’s invariant between them. For instance, since the two functions do not have ratios in common, so ratios are not meaningful in the measurement of desire. So far so good. What Zynda notes is that whenever an expected-utility representation  $(\mu, v)$  exists, then so too will another representation  $(\mu^*, v)$ , where

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

The catch is that we need adjust the operation by which  $\mu^*$  and  $v$  jointly determine  $\succsim_p$ ; this time,  $\alpha \succsim_p \beta$  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-maximisation representation*. Whenever an expected-utility representation  $(\mu, v)$  of  $\succsim_p$  exists then a valuation-maximisation representation  $(\mu^*, v)$  of  $\succsim_p$  also exists, and vice versa. Now by analogy with  $v$  and  $v^*$ , one might be tempted to infer from this fact something about *meaningfulness* in  $\mu$  and  $\mu^*$ —namely, that there’s no meaningful difference between them, that what matters is what’s invariant. As Zynda suggests,

One might point out that  $\mu^*$  is simply a linear transformation of  $\mu$ , 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* [ $\mu$  and  $\mu^*$ ]... 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, pp. 64–5, notation altered for consistency)

But there were some leaps there. While the example *does* highlight something important about the meaning of  $\mu$ , *this* is very much not it.

First note that while  $\mu$  and  $\mu^*$  share their ordinal structure, that’s not *all* they share. The transformation linking  $\mu$  and  $\mu^*$  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  $\mu_1$  and  $\mu_2$  can make a difference to preferences under an expected-utility representation, differences between ordinally-equivalent but non-identical  $\mu_1^*$  and  $\mu_2^*$  will therefore likewise matter for some preferences under the valuation-maximisation rule. The same will necessarily be true for any possible ‘redefinition’ of  $\mu$ . 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  $\mu$  to  $\mu^\dagger$ , or from  $\mu^*$  to  $\mu^\dagger$ , does not preserve difference ratios. But whenever Sally’s preferences  $\succsim_p$  can be given an expected-utility representation  $(\mu, v)$ , or a valuation-maximisation representation  $(\mu^*, v)$ , then they can also be given a *schmaluation-maximisation representation*  $(\mu^\dagger, v)$  such that  $\alpha \succsim_p \beta$  iff

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

Going further, we can even construct alternate decision-theoretic representations where not even *orderings* are preserved. For *any* transformation that takes us from  $\mu$  to some  $\mu^*$ , provided  $\mu(p) \neq \mu(q)$  iff  $\mu^*(p) \neq \mu^*(q)$ , then there will be at least one (potentially very complicated) operation by which they can be combined to generate the same preferences relative to the ‘redefined’ function  $\mu^*$  as the expected utility rule does relative to  $\mu$ . Given that, 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 trivial requirement that different degrees of belief must be assigned different values.

The problem is *meaningfulness* in the representation of any quantity is only sensibly defined relative to a fixed choice of representational format.<sup>8</sup> Ratios of numerical lengths are meaningful when lengths are additively represented in the system  $\langle \mathbb{R}^{\geq 0}, \geq, + \rangle$ , for example, but that’s only one of infinitely many numerical systems in which we might choose to represent length. Hölder (1901) showed that the system of lengths can be given a non-additive representation in the system  $\langle \mathbb{R}^{\geq 1}, \geq, \times \rangle$ , and ratios are not invariant in multiplicative representations. Nor do 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 fact that different lengths are assigned different values—not even the numerical ordering of lengths are invariant in that broader sense. Likewise for decision-theoretic representations. So the fact that ratios vary between  $\mu$  and  $\mu^*$ , or that difference ratios vary between  $\mu$  and  $\mu^\dagger$ , implies nothing whatsoever about the meaningfulness of those relations.

<sup>8</sup> More precisely, where  $\mathbf{X}$  is a set and  $R_1, R_2, \dots$  are relations defined on  $\mathbf{X}$ , suppose that  $\mathcal{Q} = \langle \mathbf{X}, R_1, R_2, \dots \rangle$  is a qualitative system and that there exists at least one structure-preserving mapping from  $\mathcal{Q}$  into a numerical system  $\mathcal{N} = \langle \mathbf{Y}, S_1, S_2, \dots \rangle$ . 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  $\varphi$  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 *Q-meaningful relative to*  $\mathcal{N}$  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  $\psi$  from  $\mathcal{Q}$  into  $\mathcal{N}$ . This amounts to saying that  $S$  is meaningful relative to a given representational format whenever it corresponds to the *same* qualitative relation regardless of how we choose to represent the qualitative system within that format. So ratios are meaningful in additive measures of length, for instance, since the qualitative meaning of those ratios doesn’t depend on an arbitrary choice of additive scale. For more discussion on meaningfulness, see (Pfanzagl 1968), (Luce 1978), (Narens 1985), and especially (Luce *et al.* 1990).

Stronger: what Zynda-style examples actually establish is that ratios in  $\mu$  really are meaningful *relative to* expected-utility representations, precisely because any transformation of  $\mu$  that does *not* preserve ratios must therefore employ a different combination 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 structure we're trying to represent is not necessarily something wholly 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 form of 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 ratios of differences in utility functions are meaningful has appealed to the role desires play as part of a broader system. Considered 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 that the need for a richer measure is manifest. Two utility functions may be ordinally-equivalent, but if they diverge in their difference ratios then they'll be differentiated in some decision situations—and therein lies the qualitative meaning of those 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.

## 5. Functional Interpretations

With so much setup, we 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  $\mu$  in  $\mathcal{R}$ . Comparativist interpretations agree on this point: where  $\mathcal{R}$  represents an incomplete confidence ordering, then every  $\mu$  in  $\mathcal{R}$  determinately misrepresents 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 our thought: if it may end up being impossible to appreciate what's *real* versus what's *artefact* in a formal model of belief without appreciating the role those models play in the theories that make use of them, then why not just take those roles themselves to be what's real? We 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  $T$  is some decision-theory-*cum*-epistemology in which representors have a role to play. In the usual functionalist manner (à la Lewis 1970), we can treat representors as theoretical terms implicitly defined by the role they play within  $T$ . According to  $T$ , the state 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 comes close to doing so. Thus we take the *meaning* of  $\mathcal{R}$  relative to  $T$  to be a function from worlds to whatever state 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 different *according to  $T$*  just in case  $\mathcal{R}$  and  $\mathcal{R}'$  play distinct roles within that theory. On a hypothetical comparativist decision theory, for instance,  $\mathcal{R}$  and  $\mathcal{R}'$  should play the same role whenever they determine the same confidence relations. Of course, whether we actually ought to treat those representors 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 some theory if we don't have much reason to think that theory is at all very plausible. Thus, we 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, pp. 204ff) once proposed something much like what we have in mind here. What they propose is that (absolute) degrees of belief are those things that play the kinds of roles numerical probabilities are supposed to play in the best systematisations of our ideas about rational belief and decision-making. They called their view *primitivism*, but they also note (2007, p. 210) 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 these 'definitions' should really be counted as such. They prefer instead to say that the concept of 'degrees of belief' is a theoretical primitive, and we get a handle on the concept by understanding the roles it plays in the theories that make use of them. It is a difference that makes no difference. The essence of Eriksson & Hájek's proposal is functionalism, broadly construed, and in that respect is closely related to ours.

But not quite the same. Eriksson & Hájek propose to take as primitive *absolute degrees of belief*. That makes sense inasmuch as we’re modelling beliefs in the traditional way, since everything a probability function says about a belief state can be derived from what it says about the absolute degree of belief it associates with each proposition. But 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 how belief state is structured as a whole and the role it plays.  $\mathcal{R}_{coin}$  assigns the very same maximally imprecise interval to  $p$  and  $q$ , but it would be a mistake to say that Sally’s attitudes towards  $p$  and  $q$  are the same. Better instead to follow Lewis: let the entire system of beliefs be our primitive, represented by the set of functions  $\mathcal{R}$ , 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.

More importantly, 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. Of course, sometimes we’ll be in a position to express differences between two representors in these terms: 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.

## References

- Alon, S. and E. Lehrer (2014). Subjective multi-prior probability: A representation of a partial likelihood relation. *Journal of Economic Theory* 151, 476–492.
- Bradley, S. (2014). Imprecise Probabilities. *Stanford Encyclopedia of Philosophy*.
- Builes, D., S. Horowitz, and M. Schoenfield (2022). Dilating and contracting arbitrarily. *Nous* 56(1), 3–20.
- Chang, R. (2002). The Possibility of Parity. *Ethics* 112(4), 659–688.
- de Finetti, B. (1931). Sul Significato Soggettivo Della Probabilità. *Fundamenta Mathematicae* 17(1), 298–329.
- Debreu, G. (1960). Topological methods in cardinal utility theory. In K. J. Arrow, S. Karlin, and P. Suppes (Eds.), *Mathematical Methods in the Social Sciences*, pp. 16–26. Stanford University Press.
- Ding, Y., W. H. Holliday, and T. F. Icard III (2021). Logics of Imprecise Comparative Probability. *International Journal of Approximate Reasoning* 132, 154–180.

- Domotor, Z. (1970). Qualitative information and entropy structures. In J. Hintikka and P. Suppes (Eds.), *Information and inference*, pp. 148–194. Reidel.
- Elliott, E. (2022a). Comparativism and the Measurement of Partial Belief. *Erkenntnis* 87, 2843–2870.
- Elliott, E. (2022b). What is ‘Real’ in Interpersonal Comparisons of Confidence. *Australasian Journal of Philosophy* 100(1), 102–116.
- Eriksson, L. and A. Hájek (2007). What are degrees of belief? *Studia Logica* 86(2), 183–213.
- Eva, B. (2019). Principles of Indifference. *The Journal of Philosophy* 116(7), 390–411.
- Eva, B. and R. Stern. Comparative Opinion Loss. *Philosophy and Phenomenological Research*.
- Fine, T. L. (1973). *Theories of Probability: An Examination of Foundations*. Academic Press.
- Fishburn, P. C. (1986). The axioms of subjective probability. *Statistical Science* 1(3), 335–345.
- Harrison-Trainor, M., W. H. Holliday, and T. F. Icard (2016). A Note on Cancellation Axioms for Comparative Probability. *Theory and Decision* 80(1).
- Harrison-Trainor, M., W. H. Holliday, and T. F. Icard III (2018). Inferring Probability Comparisons. *Mathematical Social Sciences* 91, 62–70.
- Hájek, A. (2003). What Conditional Probability Could Not Be. *Synthese* 137(3), 273–323.
- Hájek, A. and M. Smithson (2012). Rationality and indeterminate probabilities. *Synthese* 187, 33–48.
- Hölder, O. (1901). Die Axiome der Quantitat und die Lehre vom Mass. *Verh. Kgl. Sächsis. Ges. Wiss. Leipzig, Math.-Physc. Classe* 53, 1–63.
- Jeffrey, R. (1983). Bayesianism with a human face. *Testing Scientific Theories, Minnesota Studies in the Philosophy of Science* 10, 133–56.
- Joyce, J. (2005). How Probabilities Reflect Evidence. *Philosophical Perspectives* 19, 153–178.
- Joyce, J. (2010). A Defense of Imprecise Credences in Inference and Decision Making. *Philosophical Perspectives* 24(1), 281–323.
- Kahneman, D. and A. Tversky (1979). Prospect theory: An analysis of decision under risk. *Econometrica* 47(2), 263–291.
- Kaplan, M. (1996). *Decision Theory as Philosophy*. Cambridge University Press.
- Kaplan, M. (2002). Decision theory and epistemology. In P. Moser (Ed.), *The Oxford Handbook of Epistemology*, pp. 434–462. Oxford: Oxford University Press.
- Kaplan, M. (2010). In defense of modest probabilism. *Synthese* 176(1), 41–55.
- Kaplan, M. and T. Fine (1977). Joint orders in comparative probability. *The Annals of Probability* 5, 161–179.
- Keynes, J. M. (1921). *A Treatise on Probability*. New York: Macmillan.
- Konek, J. (2019). Comparative Probabilities. In *The Open Handbook of Formal Epistemology*, pp. 267–348. PhilPapers Foundation.
- Koopman, B. O. (1940a). The Axioms and Algebra of Intuitive Probability. *Annals of Mathematics* 41(2), 269–292.
- Koopman, B. O. (1940b). The Bases of Probability. *Bulletin of the American Mathematical Society* 46(10), 763–774.
- Krantz, D. H., R. D. Luce, P. Suppes, and A. Tversky (1971). *Foundations of measurement, Vol. I: Additive and polynomial representations*. Academic Press.

- Levinstein, B. (2019). Imprecise Epistemic Values and Imprecise Credences. *Australasian Journal of Philosophy* 97(4), 741–760.
- Lewis, D. (1970). How to Define Theoretical Terms. *The Journal of Philosophy* 67(13), 427–446.
- Lewis, D. (1986). *On the Plurality of Worlds*. Cambridge University Press.
- Lewis, D. (1994). Reduction of Mind. In S. Guttenplan (Ed.), *Companion to the Philosophy of Mind*, pp. 412–431. Blackwell.
- Lewis, D. (1999). Many, but almost one. In *Papers in Metaphysics and Epistemology*, pp. 164–182. Cambridge University Press.
- Luce, R. and L. Narens (1978). Qualitative independence in probability theory. *Theory and Decision* 9, 225–239.
- Luce, R. D. (1978). Dimensionally Invariant Numerical Laws Correspond to Meaningful Qualitative Relations. *Philosophy of Science* 45(1), 1–16.
- Luce, R. D., D. H. Krantz, P. Suppes, and A. Tversky (1990). *Foundations of Measurement, Vol. III: Representation, Axiomatization, and Invariance*. New York: Dover.
- Luce, R. D. and J. W. Tukey (1964). Simultaneous conjoint measurement: a new scale type of fundamental measurement. *Journal of Mathematical Psychology* 1(1), 1–27.
- Mayo-Wilson, C. and G. Wheeler (2019). Epistemic Decision Theory’s Reckoning. Manuscript. [http://philsci-archive.pitt.edu/16374/1/25a\\_EDTR.pdf](http://philsci-archive.pitt.edu/16374/1/25a_EDTR.pdf).
- Meschini, D., M. Lehto, and J. Piilonen (2004). Geometry, pregeometry and beyond. *Stud.Hist.Philos.Mod.Phys.* 36, 435–464.
- Michell, J. (2021). Representational Measurement Theory: Is its number up? *Theory & Psychology* 31, 3–23.
- Miranda, E. and S. Destercke (2015). Extreme points of the credal sets generated by comparative probabilities. *Journal of Mathematical Psychology* 64-65, 44–57.
- Narens, L. (1985). *Abstract Measurement Theory*. MIT Press.
- Nehring, K. (2009). Imprecise probabilistic beliefs as a context for decision-making under ambiguity. *Journal of Economic Theory* 144, 1054–1091.
- Pfanzagl, J. (1968). *Theory of Measurement*. New York: Wiley.
- Ramsey, F. P. (1931). Truth and probability. In R. B. Braithwaite (Ed.), *The Foundations of Mathematics and Other Logical Essays*, pp. 156–198. London: Routledge.
- Rinard, S. (2015). A Decision Theory for Imprecise Credences. *Philosopher’s Imprint* 15(7), 1–16.
- Rinard, S. (2017). Imprecise Probability and Higher Order Vagueness. *Res Philosophica* 94(2), 257–273.
- Savage, L. J. (1954). *The Foundations of Statistics*. Dover.
- Seidenfeld, T. (1988). Decision Theory Without ‘Independence’ or Without ‘Ordering’. *Economics and Philosophy* 4(2), 267–290.
- Smith, N. J. J. Interpreting Imprecise Probabilities. *Philosophical Quarterly*.
- Stefánsson, H. O. (2017). What Is ‘Real’ in Probabilism? *Australasian Journal of Philosophy* 97(3), 573–587.
- Stefánsson, H. O. (2018). On the Ratio Challenge for Comparativism. *Australasian Journal of Philosophy* 96(2), 380–390.
- Sturgeon, S. (2008). Reason and the grain of belief. *Nous* 42, 139–165.
- Tarski, A. (1954). Contributions to the theory of models I. *Indagationes Mathematicae* 16, 26–32.
- Troffaes, M. C. M. (2007). Decision making under uncertainty using imprecise probabilities. *International Journal of Approximate Reasoning* 45, 17–29.

- van Fraassen, B. (1990). Figures in a Probability Landscape. In J. Dunn and A. Gupta (Eds.), *Truth or Consequences*, pp. 345–356. Amsterdam: Kluwer.
- van Fraassen, B. (2006). Vague Expectation Value Loss. *Philosophical Studies* 127, 483–491.
- von Neumann, J. and O. Morgenstern (1944). *Theory of Games and Economic Behavior*. Princeton University Press.
- Walley, P. (1991). *Statistical Reasoning with Imprecise Probabilities*. Chapman & Hall.
- Wheeler, J. (1964). Geometrodynamics and the issue of the final state. In C. De Witt and B. S. De Witt (Eds.), *Relativity, groups and topology*, pp. 317–520. Gordon and Breach.
- Wheeler, J. A. (1980). Pregeometry: Motivations and prospects. In *Quantum theory and gravitation*. Academic Press.
- Williams, J. R. G. (2014). Decision-Making Under Indeterminacy. *Philosopher's Imprint* 14(4), 1–34.
- Zynda, L. (2000). Representation Theorems and Realism About Degrees of Belief. *Philosophy of Science* 67(1), 45–69.