The Measurement of Subjective Probability

Edward J. R. Elliott

Draft of July 31, 2023

Contents

1 Introduction **2**

2 Representation and Measurement **5**

2.1. Preliminary concepts . . . 6
2.2. Representation theorems and uniqueness . . . 9
2.3. Extensive and conjoint measurement . . . 11
2.4. Conventionality . . . 15
2.5. Meaningfulness . . . 17

3 Clarifications and Desiderata **20**

3.1. Quantitation, not elicitation . . . 21
3.2. Measurement, not metaphysics . . . 23
3.3. Simplifying assumptions . . . 26
3.4. Desiderata . . . 29

4 Epistemic Approaches: Comparative Confidence **34**

4.1. Probabilistic representations . . . 34
4.2. The problem of logical omniscience . . . 39

5 Epistemic Approaches: Alternatives **44**

5.1. Conditional confidence . . . 44
5.2. Qualitative expectations . . . 47
5.3. Multiprimitive structures . . . 50

6 Decision-Theoretic Approaches **53**

6.1. The objects of preference . . . 53
6.2. Ramsey's theorem . . . 56
6.3. Uniqueness and meaningfulness . . . 60
6.4. An apology for representation theorems . . . 63

Chapter 1

Introduction

It's a commonplace nowadays that beliefs come in degrees, though this isn't universally accepted. There are some holdouts—those who say the recent uptick of interest in "credences" and "subjective probabilities" is yet another philosophical fad that will eventually run its course. But that's hardly plausible. A very large body of work across a wide range of disciplines developed over many decades depends on the presumption that our beliefs—or something closely connected to our beliefs—admit of degrees and, moreover, that it makes good sense to represent those degrees numerically. These numerical representations of belief are far too useful for far too much to be just a passing trend.

I expect most readers will agree. But what we're much less likely to agree on is what the numbers *mean*. What is the underlying psychological reality to which these representations supposedly correspond? Perspectives on this matter vary wildly. For some, degrees of belief are understood to be explicit, on-the-fly judgements about the probability of an event, or a conscious attempt to put a number on the weight of one's evidence, or the intensity of some phenomenology of confidence experienced when contemplating a possibility. Others will, like myself, think of degrees of belief as implicit attitudes—attitudes that may be present and playing a role in your cognitive economy even if you're not consciously aware of their doing so, and even if they're not readily accessible to conscious introspection. But there's a substantial variety of perspectives, too, on how these attitudes are to be understood. If I say that Ramsey believes p to degree 0.69, does that '0.69' tell us something about p's location in Ramsey's subjective confidence ordering over possibilities? Does it tell us something about Ramsey's willingness to bet on p? About the centrality of p to Ramsey's web of belief, or his willingness to revise his opinion about p in the face of countervailing evidence? All of the above? None of the above?

A related question concerns what's *meaningful* in a numerical representation of belief. What, in other words, does it take for numerically distinct representations of belief to nevertheless represent the same thing? Most are happy to accept that there's no uniquely correct way to represent degrees of belief within a numerical framework, just as there's clearly no uniquely correct way to numerically represent lengths, or temperatures, or desirabilities. As Builes *et al.* recently put it,

> … there's nothing "0.69-ish" about my degree of confidence in p, beyond the fact that 0.69 can serve as an adequate representation of my degree of confidence within a particular representational system. But 69, for example, or 732.6 for that matter, would work just as well, provided the system was structured in the right way. (2022, 7)

But just *what it is* for the representational system to be 'structured in the right way' is about as clear as mud. And here again we find plenty of variation and disagreement. The most common numerical representations of belief use *credence functions*—mappings from propositions to real values between 0 and 1. It goes without saying that the relation induced over the propositions by their numerical ordering in a credence function is supposed to correspond to relative strengths of belief regarding those propositions. But is that *all* the meaningful information contained in a given credence function? If two credence functions are ordinally-equivalent, are they therefore equivalent in meaning? If so, then we'd better get to work revising the many theories of rational belief and decision-making that presuppose meaningful differences between ordinally-equivalent credence functions. On the other hand, if there's more to meaning than numerical orderings, then just what additional structure *is* relevant?

And even where we find agreement regarding *what* is meaningful, there's often still plenty of disagreement on the *how* and the *when*. According to one approach to these sorts of questions—what I'll be calling the *epistemic approach*—a system of beliefs admits of numerical representation just in case it has a certain kind of internal structure that can be mirrored in an appropriate numerical framework. A rather different tack—the *decision-theoretic approach*—focuses instead on the relationship between beliefs, desires, and preferences in the context of decision-making. Roughly: is the numerical representability of belief a matter of having a coherent system of beliefs, or a matter of having beliefs that relate to preferences and desires in a coherent way? These can end up saying much the same sorts of things about what should and should not be considered meaningful in our numerical representations of belief, but they can diverge quite significantly on the matter of what an agent must be like in order for their beliefs to admit of such representations in the first place.

This *Element* is about the measurement of belief. In sum: what do our numerical representations of belief actually represent, how do they represent it, and under what conditions are such representations meaningful? Cards on the table: I prefer the decision-theoretic approach.[1] But I'll not spend a great deal of time arguing in favour of my own view, nor arguing against the competitors. I mean—I'll do a little of that here and there, and my biases will be apparent in parts of the discussion, but the main purpose of this work is expositional rather than

[1] More carefully, I think the decision-theoretic approach supplies the most plausible way to interpret numerical representations of belief in the Bayesian tradition, especially in decision-theoretic contexts but also in the context of much if not most Bayesian epistemology. Bayesian theories and arguments in that tradition routinely make assumptions about meaningfulness that are harder to capture with an epistemic approach. But it would be grossly implausible to assume that there's only one way to understand the measurement of belief, and for some theoretical applications some epistemic approaches are clearly adequate.

argumentative. So I'll focus more on explaining what the epistemic and decision-theoretic approaches are, some of the variation within those two approaches, and the implications they have regarding what kinds of numerical representations are possible, when they're possible, and what ought to be considered meaningful in those representations.

The remainder of the discussion proceeds as follows. Chapter 2 introduces some key concepts from the theory of measurement, while Chapter 3 provides some clarifications and general assumptions/desiderata regarding a theory of belief measurement. We then turn to the epistemic approaches: Chapter 4 covers the simplest version of the epistemic approach, built around binary comparative confidence relations, while Chapter 5 gives an overview of several alternatives. Finally, Chapter 6 gives and overview of the decision-theoretic approach, discusses one version in some detail, and addresses some common misunderstandings and objections.

Chapter 2

Representation and Measurement

We find it abundantly useful to express many physical facts using numbers and numerical relations. There's no great mystery to this, even for the mathematical Platonist who thinks that numbers and numerical relations are abstracta and not present in the physical world in the same manner as electrons or chairs or gravitational attraction. When I say that I've gained *at least 2 pounds* thanks to all the nice food at recent conference, which is *twice as much* as what I gained at the last conference, I'm using those numbers and numerical relations to refer to and reason about my ever-increasing weight. These claims aren't made true by virtue of any little numbers attached somewhere to my body, slowly and inevitably going up over time. Rather, the numbers and numerical relations serve as abstract stand-ins for physical properties and physical relations, and they do this by virtue of structural similarity between them.

What we call *quantities* are determinable properties whose determinates have a certain kind of salient relational structure rendering them ripe for numerical representation. Length, for instance, is a determinable attribute, with determinates—the specific lengths—sharing higher-order relations between them that can be usefully represented within a numerical framework. For any two physical objects o and o' and a fixed orientation for each, either o will be at least as long as o', or o' will be at least as long as o, or both. Here, the *at least as long* holds between physical objects; but we can also understand it as a second-order relation between the length attributes directly. Say that any two objects have the same length, L, if each is at least as long as the other. Say next that L is *at least as long* as L' just in case any object with property L is at least as long as any object with property L'. We can then associate any two lengths L and L' with real numbers x and y in such a manner that L is at least as long as L' just in case $x \geq y$.

In this example, the lengths L, L' and the *at least as long* relation between them are said to be *qualitative*, whereas the numbers x, y and the $\geq$ relation between them serve as their *numerical* representations. Think of a qualitative property or relation as one that can be characterised without explicit reference to numbers or numerical relations. So 'qualitative' here contrasts with 'numerical', not with 'quantitative'—the idea being that quantities can be characterised in either qualitative terms or numerical terms, with the latter being possible precisely because the abstract numerical stuff shares a structure in common with the real-world qualitative stuff it represents.[1]

The purpose of this chapter is to expand on that initial idea and make it more precise; more generally, to introduce some key concepts for discussing the numerical representation of quantities. I start with the fundamentals of the Representational Theory of Measurement (RTM).[2]

2.1 Preliminary concepts

I presume familiarity with predicate logic, and with the elementary concepts and notation of set theory. Much of what follows will revolve around binary relations and operations, though, so the following are worth highlighting:

Definition 2.1 An n-ary relation on a set $\mathbf{X}$ is a subset of $\mathbf{X}^n$. Where $R \subseteq \mathbf{X} \times \mathbf{X}$, by convention, xRy iff $(x, y) \in R$ and $x \not R y$ iff $(x, y) \notin R$; furthermore, we say R is

- *transitive* iff, for all $x, y, z \in \mathbf{X}$, xRy and yRz implies xRz
- *complete* iff, for all $x, y \in \mathbf{X}$, xRy or yRx
- *reflexive* iff, for all $x \in \mathbf{X}$, xRx
- *symmetric* iff, for all $x, y \in \mathbf{X}$, xRy implies yRx
- *asymmetric* iff, for all $x, y \in \mathbf{X}$, xRy implies not yRx
- *antisymmetric* iff, for all $x, y \in \mathbf{X}$, xRy and yRx implies $x = y$
- an *equivalence relation* iff R is transitive, reflexive and symmetric
- a *preorder* iff R is transitive and reflexive
- a *weak order* iff R is a complete preorder
- a *total order* iff R is an antisymmetric complete preorder

Preorders—especially weak orders—will be important. Throughout, I'll use $\succsim$ to represent a variety of qualitative preorder relations, and I'll use $\sim$ and $\succ$ for the symmetric and asymmetric parts of $\succsim$ respectively. That is, and unless otherwise specified, I'll take it as read that:

- $x \sim y$ iff $x \succsim y$ and $y \succsim x$
- $x \succ y$ iff $x \succsim y$ and $y \not\succsim x$

Furthermore, where $\succsim$ is defined on a set $\mathbf{X}$, then x is said to be:

- *minimal* (in $\succsim$) iff $y \succsim x$ for all $y \in \mathbf{X}$
- *maximal* (in $\succsim$) iff $x \succsim y$ for all $y \in \mathbf{X}$

[1] This usage of 'qualitative' is common in the RTM literature. Some will say that a numerical system is defined by its structure; hence anything with the same structure instantiates that system and should also be considered 'numerical' (e.g., Michell 2021). That may be right. But what I have to say won't hinge on whether 'qualitative' systems *instantiate* 'numerical' systems or are *represented by* them, and in either case the terminological distinction proves useful.

[2] The *locus classicus* for the RTM is (Krantz et al. 1971); see also (Suppes & Zinnes 1963), (Pfanzagl 1968), (Narens 1985), and (Roberts 1985).

Definition 2.2 An n-ary operation on a set $\mathbf{X}$ is a (total or partial) function from $\mathbf{X}^n$ into $\mathbf{X}$. Suppose that $\bullet : \mathbf{X}\times\mathbf{X} \mapsto \mathbf{X}$. By convention, $x \bullet y = z$ iff $\bullet(x, y) = z$, and $x \bullet y$ is *defined* iff $\bullet(x, y)$ is defined. Furthermore, we say $\bullet$ is

- *total* iff $x \bullet y$ is defined for all $x, y \in \mathbf{X}$, otherwise *partial*
- *commutative* iff $\bullet$ is total and for all $x, y \in \mathbf{X}$, $x \bullet y = y \bullet x$
- *associative* iff $\bullet$ is total and for all $x, y \in \mathbf{X}$, $x \bullet (y \bullet z) = (x \bullet y) \bullet z$

Note that *properties* are just the special case of n-ary relations where $n = 1$, and that every n-ary operation can be recast as an $(n+1)$-ary relation. For example, addition is a total binary operation on the set of real numbers $\mathbb{R}$, since it maps $\mathbb{R}\times\mathbb{R}$ back into $\mathbb{R}$; it can also be construed as the ternary relation R on $\mathbb{R}$ such that $(x, y, z) \in R$ iff $x + y = z$. As such, I'll usually just write 'relations' rather than 'properties and relations', and I'll sometimes use 'relations' to also cover operations—but where I intend to refer to operations in particular this will be explicitly marked.

Next we need the generic notion of a *relational system*. This is a system comprising a set, one or more distinguished relations on that set, and zero or more distinguished binary operations:

Definition 2.3 Let $\mathbf{I}$ $(\supset \varnothing)$ and $\mathbf{J}$ $(\supseteq \varnothing)$ be index sets. Then $\langle \mathbf{X}, R_i; \bullet_j \rangle\,^{i\in\mathbf{I}}_{j\in\mathbf{J}}$ is a *relational system* iff $\mathbf{X}$ is a non-empty set, the R_i are relations on $\mathbf{X}$, and the $\bullet_j$ are binary operations on $\mathbf{X}$.

The relations and operations used to characterise a relational system are known as the *primitives* of that system. Note the use of the semi-colon to explicitly separate the primitive relations from the primitive operations.[3]

An example of a simple relational system is $\langle \mathbb{R}, \geq \rangle$, comprising the set of reals $\mathbb{R}$ and the primitive relation $\geq$ on $\mathbb{R}$. A richer relational system would be $\langle \mathbb{R}, \geq; + \rangle$, which includes also the primitive binary operation $+$. These are what we'll call *numerical systems*—they're comprised of a set of numbers and one or more relations thereupon. More generally, we take a numerical system to be any relational system constructed from numerical stuff. There's no need for us to be very precise about 'numerical stuff'—some relational systems have a numerical feel about them, and that'll suffice for referring to them as numerical systems. In contrast there are *qualitative systems*, or systems constructed from qualitative stuff. For example, if $\mathbf{L}$ is the set of determinate length properties as characterised at the beginning of the chapter, and $\succsim$ is the *at least as long* relation, then $\langle \mathbf{L}, \succsim \rangle$ will count as a qualitative system. Likewise, for any two lengths L and L', we let their *end-to-end concatenation*, $L \circ L'$, be the length L'' of any object that's as long as what you get when you take two disjoint objects of length L and L' and joint them end-to-end. (See Figure 2.1.) Then $\circ$ will be an operation on $\mathbf{L}$, and $\langle \mathbf{L}, \succsim; \circ \rangle$ will be a qualitative system.

[3] I follow Roberts (1985) in how I define relational systems. Doing so allows for the distinction between weak homomorphisms and strong homomorphisms (Defintion 2.4), which helps resolve some minor technical problems that arise in connection to the representation of partial operations and non-antisymmetric preorders.

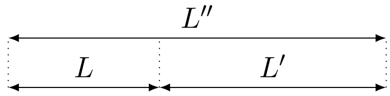

Figure 2.1: L'' is the end-to-end concatenation of L and L' ($L \circ L' = L''$)

Henceforth I'll use $\mathcal{N}$ for numerical systems and $\mathcal{Q}$ for qualitative systems. Next we need to a way of expressing when a numerical relational system possesses a similar structure to that of a qualitative system, such that the former might be employed to represent the latter. We can do this via the notion of a *structure-preserving mapping*, or *homomorphism*:

Definition 2.4 Let $\mathcal{Q} = \langle \mathbf{X}, R_i; \bullet_j \rangle$ and $\mathcal{N} = \langle \mathbf{Y}, S_i; *_j \rangle$, where $i \in \mathbf{I}$ and $j \in \mathbf{J}$. Then $\varphi : \mathbf{X} \mapsto \mathbf{Y}$ is a *weak homomorphism from $\mathcal{Q}$ into $\mathcal{N}$* iff

1. R_i is an n-ary relation iff S_i is an n-ary relation
2. $(x_1, \ldots, x_n) \in R_i$ iff $(\varphi(x_1), \ldots, \varphi(x_n)) \in S_i$
3. $\varphi(x \bullet_j y) = \varphi(x) *_j \varphi(y)$

φ is a *strong homomorphism from $\mathcal{Q}$ into $\mathcal{N}$* if, in addition,

4. $x \bullet_j y = z$ iff $\varphi(x) *_j \varphi(y) = \varphi(z)$

Corresponding to the distinction between weak homomorphisms and homomorphisms, we can say that φ *weakly maps* the operation $\bullet$ into the operation $*$ whenever

$$x \bullet y = z \text{ implies } \varphi(x) * \varphi(y) = \varphi(z)$$

and *strongly maps* $\bullet$ into $*$ whenever the converse also holds. An example will help to make all this clearer. Start first with the simpler qualitative system $\langle \mathbf{L}, \succsim \rangle$. A function $\varphi : \mathbf{L} \mapsto \mathbb{R}$ is a homomorphism from $\langle \mathbf{L}, \succsim \rangle$ into $\langle \mathbb{R}, \geq \rangle$ when:

$$L \succsim L' \text{ iff } \varphi(L) \geq \varphi(L')$$

Note, of course, that since there are no primitive operations in $\langle \mathbf{L}, \succsim \rangle$, conditions 3 and 4 are trivially satisfied. (In this case we don't bother with the weak/strong distinction.) Next, consider the richer system $\langle \mathbf{L}, \succsim; \circ \rangle$, this time endowed with a primitive concatenation operation. A function $\varphi : \mathbf{L} \mapsto \mathbb{R}$ is a weak homomorphism from $\langle \mathbf{L}, \succsim; \circ \rangle$ into $\langle \mathbb{R}, \geq; + \rangle$ whenever, in addition to the above, it weakly maps $\circ$ into $+$:

$$\varphi(L \circ L') = \varphi(L) + \varphi(L')$$

And φ is a strong homomorphism if it strongly maps $\circ$ into $+$:

$$\varphi(L'') = \varphi(L) + \varphi(L') \text{ iff } L'' = L \circ L$$

If $\circ$ is a total operation and $\succsim$ is antisymmetric then every weak homomorphism from $\langle \mathbf{L}, \succsim; \circ \rangle$ into $\langle \mathbb{R}, \geq; + \rangle$ will be a strong homomorphism, but otherwise this needn't be the case.

2.2 Representation theorems and uniqueness

A homomorphism maps the primitive relations and operations of one system into the primitive relations and operations of another. When a weak homomorphism from $\mathcal{Q}$ into $\mathcal{N}$ exists, we can say that $\mathcal{N}$ has—or otherwise includes as a proper part—a formal structure that's similar to that of $\mathcal{Q}$. A strong homomorphism establishes a stronger similarity of structure. In either case, it's is this similarity of structure that allows us to represent $\mathcal{Q}$ using (or 'in') $\mathcal{N}$. Thus, the central theoretical objects of the RTM are theorems that establish conditions for when a qualitative system $\mathcal{Q}$ can be represented in some specific numerical system $\mathcal{N}$. These are known as *representation theorems*.

Henceforth, let $\mathbf{\Phi}(\mathcal{Q}, \mathcal{N})$ denote the set of all weak homomorphisms from $\mathcal{Q}$ into $\mathcal{N}$. Then, for a prespecified $\mathcal{N}$, a representation theorem will provide sufficient conditions on $\mathcal{Q}$ to guarantee $\mathbf{\Phi}(\mathcal{Q}, \mathcal{N}) \neq \varnothing$. The conditions are usually called the *axioms* of that theorem. Typically, the axioms will be chosen so that most of them are necessary for representability. That is, they are consequences of the assumption that some homomorphism exists. A representation theorem will also typically include at least a few axioms that aren't necessary for representability. These are usually known as *structural* axioms.[4] For example:

Theorem 2.1 (Krantz et al. 1971, 15) *Let* $\mathbf{X}$ *be a set and* $\succsim$ *a binary relation on* $\mathbf{X}$*. Then there is at least one homomorphism from* $\langle \mathbf{X}, \succsim \rangle$ *into* $\langle \mathbb{R}, \geq \rangle$ *if*

1. $\mathbf{X}$ *is finite* *(finitude)*
2. $\succsim$ *is a weak order* *(weak order)*

The *weak order* axiom is necessary: since $\geq$ is a weak order on $\mathbb{R}$ and $\mathbf{X}$ is mapped into $\mathbb{R}$, $\succsim$ must itself be a weak order if it's mapped into $\geq$. By contrast, *finitude* is structural—it's entirely possible to represent $\langle \mathbf{X}, \succsim \rangle$ in $\langle \mathbb{R}, \geq \rangle$ even if $\mathbf{X}$ is infinite, though in that case additional axioms are needed to ensure representability. (See Krantz *et al.* 1971, 40–1, for details.)

A representation theorem will also usually include some form of *uniqueness result*. In the ideal case, the uniqueness result tells us about the relationship between homomorphisms belonging to $\mathbf{\Phi}(\mathcal{Q}, \mathcal{N})$ for all $\mathcal{Q}$ satisfying the axioms of the associated representation theorem. Continuing the example, it's plain to see that if φ is any homomorphism from $\langle \mathbf{X}, \succsim \rangle$ into $\langle \mathbb{R}, \geq \rangle$, then so too is $\psi : \mathbf{X} \mapsto \mathbb{R}$ iff

$$\psi(x) \geq \psi(y) \text{ iff } \varphi(x) \geq \varphi(y)$$

Any ψ satisfying this condition is related to φ by a *strictly increasing* (or *order-preserving*) transformation. Given that, the kind of uniqueness result we'd expect to find attached to Theorem 2.1 would say that given *weak order* and *finitude*, the φ in $\mathbf{\Phi}(\langle \mathbf{X}, \succsim \rangle, \langle \mathbb{R}, \geq \rangle)$ are unique up to an order-preserving transformation. The 'unique up to' phrasing can be read as saying that the homomorphism set is constrained by the specified form of transformation; it designates a property shared by all and only the functions in the set.

[4] Be careful: an axiom might be necessary for a representation theorem, but not necessary for representability. This is because a representation theorem often do more than simply assert sufficient conditions for representability.

Two points of caution about uniqueness results. First: a uniqueness result applies to all systems *satisfying the axioms of the associated representation theorem*, and not necessarily to all systems that are representable in the numerical system $\mathcal{N}$. This is important if the representation theorem includes structural axioms, which are sometimes included to strengthen the uniqueness result. As a rule of thumb, the more structural constraints imposed on $\mathcal{Q}$, the more restricted the potential homomorphisms from $\mathcal{Q}$ into $\mathcal{N}$, and hence the stronger the uniqueness result. Second: many uniqueness results only apply to a proper subset of the possible homomorphisms in $\mathbf{\Phi}(\mathcal{Q}, \mathcal{N})$. For example, the result may assert that there is only one homomorphism from $\langle \mathbf{X}, \succsim \rangle$ into $\langle \mathbb{R}, \geq \rangle$ that satisfies such-and-such properties (e.g., *being a probability measure*), even while there are infinitely many homomorphisms that do not. For these reasons, one must be careful when interpreting a uniqueness result—some results that on first glance appear rather impressive may end up only really reflecting the strength of the structural conditions employed in the representation theorem and/or arbitrary restrictions to a particular subset of homomorphisms.

Moving on—the final thing to do in this section is outline the different *scale types*. In the example above, the φ in $\mathbf{\Phi}(\mathcal{Q}, \mathcal{N})$ are unique up to order-preserving transformations. In that case, $\mathbf{\Phi}(\mathcal{Q}, \mathcal{N})$ is said to be an *ordinal scale* of $\mathcal{Q}$, and the φ in $\mathbf{\Phi}(\mathcal{Q}, \mathcal{N})$ are also called *ordinal scales* of $\mathcal{Q}$. (The ambiguity is unfortunate but well-entrenched, and context usually suffices for disambiguation.) Three other scale types will be important. The next is an *interval scale*: $\mathbf{\Phi}(\mathcal{Q}, \mathcal{N})$ is an interval scale when the $\varphi \in \mathbf{\Phi}(\mathcal{Q}, \mathcal{N})$ are unique up to a *positive affine* (or *interval-preserving*) transformation—i.e., if $\varphi \in \mathbf{\Phi}(\mathcal{Q}, \mathcal{N})$ then so is ψ, for any ψ defined such that for some real values r and s, with $r > 0$,

$$\psi(x) = r\varphi(x) + s$$

Whereas order-preserving transformations merely preserve orderings, interval-preserving transformations preserve ratios of differences (and thus also orderings). If φ and ψ are related by an interval-preserving transformation, then

$$\frac{\varphi(x) - \varphi(y)}{\varphi(z) - \varphi(w)} = \frac{\psi(x) - \psi(y)}{\psi(z) - \psi(w)}$$

Next are *ratio scales*: $\mathbf{\Phi}(\mathcal{Q}, \mathcal{N})$ is a ratio scale when the $\varphi \in \mathbf{\Phi}(\mathcal{Q}, \mathcal{N})$ are unique up to a *positive similarity* (or *ratio-preserving*) transformation—i.e., if φ is in $\mathbf{\Phi}(\mathcal{Q}, \mathcal{N})$ then so is ψ, for any ψ defined such that for some real value $r > 0$,

$$\psi(x) = r\varphi(x)$$

Ratio-preserving transformations preserve ratios (and thus also ratios of differences). If φ and ψ are related by a ratio-preserving transformation, then

$$\frac{\varphi(x)}{\varphi(y)} = \frac{\psi(x)}{\psi(y)}$$

Finally, there are *absolute scales*. This corresponds to the case where $\mathbf{\Phi}(\mathcal{Q}, \mathcal{N})$ contains exactly one homomorphism.

scale type	uniqueness condition	relations preserved
ordinal	strictly increasing transformations	orderings
interval	positive affine transformations	difference ratios
ratio	positive similarity transformations	ratios
absolute	identity transformation	everything

The foregoing classification scheme originates with (Stevens 1946). It's the most widely-known means of classifying scale types by a wide margin. It works well for most purposes, and it'll suffice for ours, though it's not the only classification scheme nor is it the most general. (A more general classification scheme, though also more complicated, can be found in Narens 1981.)

2.3 Extensive and conjoint measurement

Of special interest are 'additive' representations. Roughly, these are representations that make use of addition in some important way. This can be a little hard to define precisely, though, as what it takes for a representation to count as 'additive' can vary across measurement structures. The simplest case is that of *extensive measurement*. Here, we can say that a homomorphism from $\mathcal{Q}$ into $\mathcal{N}$ is *strongly additive* when it strongly maps one of $\mathcal{Q}$'s primitives into addition; *weak additivity* can then be defined in the obvious parallel way. The qualitative operation that gets mapped into addition is usually referred to as a *concatenation* operation.

Let's consider one example of an extensive measurement structure in a bit more detail—a *positive concatenation structure*. Since it doesn't make sense to speak of lengths shorter than no length at all, we conventionally measure length using strongly additive homomorphisms from $\langle \mathbf{L}, \succsim; \circ \rangle$ into $\langle \mathbb{R}^{\geq 0}, \geq; + \rangle$, where $\mathbb{R}^{\geq 0}$ is the set of real numbers not smaller than zero. The *meter* scale is one such homomorphism. Let L_m be the *meter length*, defined as the length of the path that light travels in a vacuum in one $299,792,458^{\text{th}}$ of a second. Then the meter scale, m, corresponds to the unique strong homomorphism from $\langle \mathbf{L}, \succsim; \circ \rangle$ into $\langle \mathbb{R}^{\geq 0}, \geq; + \rangle$ that assigns the unit value to L_m. In other words,

a) $m(L) \geq 0$ and $m(L_m) = 1$
b) $L \geq L'$ iff $m(L) \geq m(L')$
c) $L \circ L' = L''$ iff $m(L) + m(L') = m(L'')$

This method of measuring length is possible precisely because the behaviour of $\succsim$ and $\circ$ is mirrored by the behaviour of $+$ and $\geq$ over the non-negative reals. The most important conditions are:

1. $\succsim$ is a weak order (*weak order*)
2. $L \circ (L' \circ L'') = (L \circ L') \circ L''$ (*associativity*)
3. $L \circ L' = L' \circ L$ (*commutativity*)
4. $L \succsim L'$ iff $L \circ L'' \succsim L' \circ L''$ (*monotonicity*)
5. $L \circ L' \succsim L$ (*weak positivity*)
6. $L \circ L' \sim L$ only if L' is minimal in $\succsim$ (*minimal identity*)

Compare, for $x, y, z \in \mathbb{R}^{\geq 0}$:

1. $\geq$ is a weak order (*weak order*)
2. $x + (y + z) = (x + y) + z$ (*associativity*)
3. $x + y = y + x$ (*commutativity*)
4. $x \geq y$ iff $x + z \geq y + z$ (*monotonicity*)
5. $x + y \geq x$ (*weak positivity*)
6. $x + y = x$ only if $y = 0$ (*minimal identity*)

Epistemic approaches to the measurement of belief focus on representing the internal structure of the belief system, and typically posit measurement structures that look a great deal like this. It's important to note, though, that not all 'additive' representations follow the same model—they do not all require a primitive concatenation operation that gets mapped into addition. An alternative way to generate 'additive' representations employs *conjoint measurement structures*, wherein multiple quantities are represented simultaneously and the additive structure of the representation is derived from the nature of their law-like relationships. Since conjoint measurement is important for decision-theoretic approaches to the measurement of belief, it's worth considering an example in a bit of detail. The procedure is somewhat more complicated than the case of extensive measurement (see Figure 2.2).[5]

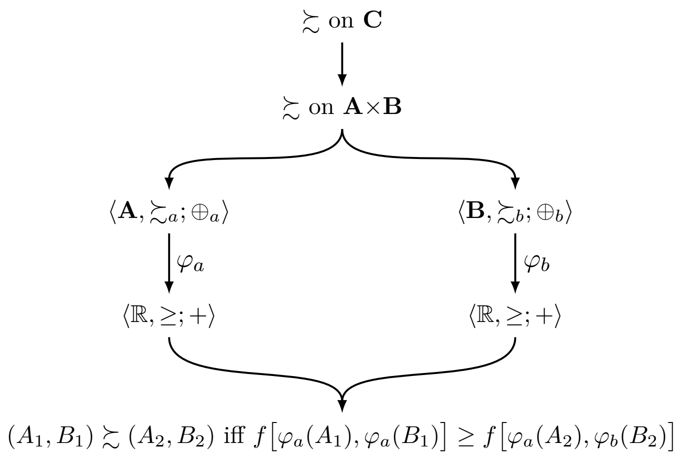

Figure 2.2: conjoint measurement structure

We start with a single weak ordering, $\succsim$, defined for some quantity $\mathbf{C}$ that's determined by two independent factors $\mathbf{A}$ and $\mathbf{B}$. For example, suppose that $\mathbf{C}$ is discomfort as determined by temperature $\mathbf{A}$ and humidity $\mathbf{B}$ (Krantz *et al.* 1971, 17–18), momentum as determined by mass and velocity (Luce & Tukey 1964, 4–5), or overall value as determined by monetary and sentimental value.

[5] The example is chosen to highlight a few key ideas; it's far from the only way conjoint measurement might proceed and different in some respects than typical decision-theoretic structures. As with extensive measurement, there's a wide variety of conjoint measurement structures and a correspondingly wide variety of numerical systems in which they can be represented.

In any case, we suppose that $\succsim$ on **C** is determined by these two factors **A** and **B**. Formally we can represent this by reconstructing $\succsim$ as an ordering not over **C** directly but instead over **A**×**B**. So, for example,

$$(A_1, B_1) \succsim (A_2, B_2)$$

is interpreted to mean that the level of **C** determined by the combination of A_1 and B_1 is at least as great as the level of **C** determined by A_2 and B_2, where A_1, A_2 are levels of **A** and B_1, B_2 are levels of **B**.

The next step is to extract from $\succsim$ two extensive 'subsystems' for **A** and **B** separately. We start by defining an ordering $\succsim_a$ over **A** by comparing the levels of **C** that result from varying the levels of **A** while holding the level of **B** fixed. That is,

$$A_1 \succsim_a A_2 \text{ iff } (A_1, B_i) \succsim (A_2, B_i) \text{ for all } B_i \in \mathbf{B}$$

So A_1 is greater than A_2 when A_1 contributes more to **C** than A_2 does, holding the level of **B** fixed. Note, of course, that to ensure $\succsim_a$ is a weak order, we need suppose that if changing from A_1 to A_2 increases the level of **C** while holding the level of **B** fixed for any particular level of **B**, then the same should hold for all levels of **B**. Essentially this amounts to saying that the contribution the **A** factor makes to **C** is independent of the contribution made by the **B** factor. This is established by the *independence* axiom noted below. An exactly parallel process gets us an ordering $\succsim_b$ over **B**.

At this point we've got two very simple subsystems, $\langle \mathbf{A}, \succsim_a \rangle$ and $\langle \mathbf{B}, \succsim_b \rangle$, and axioms appropriate for an ordinal-scale representation. But we should like extensive structures so as to enable a richer numerical representation. Thus we will need to construct a concatenation operation. Assume that **A** and **B** combine in an 'additive' fashion. (This will be qualitatively expressed by means of the *independence* and *double cancellation* axioms below.) Then, it will be possible to draw meaningful correlations in size between intervals in $\succsim_a$ and in $\succsim_b$ by comparing the effects on the level of **C** that result from varying one factor while holding the other fixed. For suppose there are A_1, A_2, B_1, B_2 such that

$$(A_1, B_2) \sim (A_2, B_1) \succ (A_1, B_1),$$

We can read this as saying that increasing from A_1 to A_2 (while holding the **B**-level fixed) has the same effect on **C** as increasing from B_1 to B_2 (while holding the **A**-level fixed). If we let $A_i \rightarrow A_j$ designate the interval between A_i and A_j as observed in the effect on **C**, and likewise for $B_i \rightarrow B_j$ *mutatis mutandis*, then what we've said is that $A_1 \rightarrow A_2$ is equal to $B_1 \rightarrow B_2$, and thus we compare intervals in one factor to intervals in the other.

Furthermore, if there are minimal levels A_0 and B_0 of **A** and **B**, we can then even go on to define separate concatenation operations $\oplus_a$ and $\oplus_b$ for each of **A** and **B**. Starting with $\oplus_a$, we say

$$A_1 \oplus_a A_2 = A_3$$

just in case the effect on **C** that results from increasing A_0 to A_3 while holding the level of **B** fixed at B_0 is equal to the effect on **C** that results from increasing the level of $\mathbf{A}$ from A_0 to A_1 and increasing the level of $\mathbf{B}$ from B_0 to some level B_x such that the result is equal in effect on $\mathbf{C}$ as observed from an increase from A_0 to A_2. That is, if

$$(A_3, B_0) \sim (A_1, B_x),$$

then $A_0 \rightarrow A_3$ is equal to $A_0 \rightarrow A_1$ plus $B_0 \rightarrow B_x$, where the latter is equal to $A_0 \rightarrow A_2$. Letting A_0 and B_0 be 'zero' points, the 'size' of the interval $A_0 \rightarrow A_i$ gives the 'size' of A_i, so this essentially amounts to saying that A_3 equals A_1 plus A_2.

Given the appropriate axioms on $\succsim$, then, we can extract subsystems $\langle \mathbf{A}, \succsim_a ; \oplus_a \rangle$ and $\langle \mathbf{B}, \succsim_b ; \oplus_a \rangle$ out of the initial system $\langle \mathbf{A} \times \mathbf{B}, \succsim \rangle$, which will admit of separate additive representations φ_a and φ_b. The final step is to show that there is a numerical operation, f, that combines φ_a and φ_b so as to represent $\succsim$ on $\mathbf{A} \times \mathbf{B}$; i.e.,

$$(A_1, B_1) \succsim (A_2, B_2) \text{ iff } f[\varphi_a(A_1), \varphi_b(B_1)] \geq f[\varphi_a(A_2), \varphi_b(B_2)]$$

The function f may take a wide variety of forms depending on the shape of $\succsim$, but one very simple case is when φ_a and φ_b combine additively to determine a final value that represents the overall level of $\mathbf{C}$:

$$(A_1, B_1) \succsim (A_2, B_2) \text{ iff } \varphi_a(A_1) + \varphi_a(B_1) \geq \varphi_a(A_2) + \varphi_b(B_2)$$

The end result is a *conjoint* representation of all three quantities $\mathbf{A}$, $\mathbf{B}$ and $\mathbf{C}$ simultaneously, achieved via a two-component vector homomorphism φ from $\langle \mathbf{A} \times \mathbf{B}, \succsim \rangle$ into $\langle \mathbb{R} \times \mathbb{R}, \geq \rangle$ that 'decomposes' into φ_a and φ_b via f.

All of this of course requires that $\succsim$ satisfies the axioms required for existence of such a representation. These axioms will essentially assert that $\succsim$ behaves in the manner expected if levels of $\mathbf{C}$ were determined by the sum of two independent factors $\mathbf{A}$ and $\mathbf{B}$. For instance, very typical necessary axioms for additive conjoint measurement structures will be:

1. $\succsim$ is a weak order (*weak order*)
2. For all $A_i, A_j, A_k, A_l \in \mathbf{A}$ and $B_i, B_j, B_k, B_l \in \mathbf{B}$, $(A_i, B_k) \succsim (A_j, B_k)$ iff $(A_i, B_l) \succsim (A_j, B_l)$, and $(A_k, B_i) \succsim (A_k, B_j)$ iff $(A_l, B_i) \succsim (A_l, B_j)$ (*independence*)
3. For all $A_i, A_j, A_k \in \mathbf{A}$ and $B_i, B_j, B_k \in \mathbf{B}$, $(A_i, B_j) \succsim (A_j, B_k)$ and $(A_j, B_i) \succsim (A_k, B_j)$ implies $(A_i, B_i) \succsim (A_k, B_k)$ (*double cancellation*)

Again, it's helpful to compare with the intended numerical representation. The *independence* axiom is straightforward:

$$\begin{gathered} x + z \geq y + z \text{ for some } z \\ \downarrow \\ x + z \geq y + z \text{ for all } z \end{gathered}$$

The *double cancellation* axiom is a little less obvious; it concerns cases in which the common terms of two inequalities cancel out to determine a third:

$$x + m \geq y + o$$
$$y + n \geq z + m$$
$$\downarrow$$
$$x + n \geq z + o$$

Let's sum up. In the example just outlined, the numerical representations of **A**, **B**, and **C** come as a package deal. Or, more accurately, they're three parts of a single representational system comprising several functions and a rule that ties them together. Note in particular—and this will be important—that the two constructed subsystems only make sense as parts of the larger system. The primitives of $\langle \mathbf{A}, \succsim_a; \oplus_a \rangle$, for instance, are characterised in terms of how **A** relates to **B** in the determination of **C**, and the operation $\oplus_a$ needn't correspond to any 'natural' concatenation operation that can be defined in terms of **A** alone. To the extent that **A** is represented as having an 'additive' structure, then, that structure is manifest in its connection to **B** and **C**. Thus, the *meaning* of the representation φ_a in this context can only be fully grasped by reference to its relation to φ_b as specified by the rule f by which they combine to represent **C**. The three numerical representations are, in that sense, inseparable.

Contrast this with the extensive measurement of $\langle \mathbf{L}, \succsim; \circ \rangle$, where the primitives of that system are characterised without any necessary reference to other quantities. One can understand what it is for the system of lengths to have an 'additive' structure just by considering how length attributes relate to other determinate length attributes—one needn't embed the system of lengths into a larger relational structure involving multiple quantities in order to comprehend what it is for one length to be *twice as long* as another, for example. An intuitive way to characterise the difference, then, is to say extensive measurement is geared towards representing the internal relational structure of a single determinable attribute, whereas the conjoint measurement is geared more towards representing the relationships that hold between several attributes.

2.4 Conventionality

One of the more important lessons of the RTM concerns the extent to which our use of numbers to represent the world is grounded in convention. By 'conventional', I mean unforced from a purely mathematical point of view, and setting aside pragmatic considerations. Most will be familiar already with some degree of conventionality in this sense, as for instance when we conventionally choose between *meters* or *inches* or *football fields* to measure length. This is conventionality in the *choice of scale*—that is, in the choice of a specific homomorphism from $\mathbf{\Phi}(\mathcal{Q}, \mathcal{N})$, given a choice of $\mathcal{Q}$ and $\mathcal{N}$.

A deeper and not as widely appreciated form of conventionality arises in the *choice of numerical system.* A very simple example is the choice to use $\geq$ to represent a qualitative weak order $\succsim$, rather than $\leq$. Either would obviously work just as well as the other. But a more complicated example is also worth mentioning. As I noted in the previous section, conventional measures of length are almost always additive homomorphisms from $\langle \mathbf{L}, \succsim; \circ \rangle$ to $\langle \mathbb{R}, \geq; + \rangle$. On any such measure, the value assigned to $L \circ L$ will always be *twice* the value assigned to L, and our overwhelming familiarity with these measures can lead to the sense that there's something uniquely correct about this—that the qualitative relation that holds between L and $L \circ L$ is an essentially *twice*-ish relation. As Brian Ellis (1968, 83) put it, there's a common sense that the 'twice' in 'twice as long' has a significance independent of the conventions of measurement. Clearly, 2 meters is *twice as long* as long as 1 meter, and 4 meters is *twice as long* as 2 meters, and so on—that is the natural and obvious way to describe these relations.

However, the axioms that permit extensive measurement—*associativity*, *commutativity*, *monotonicity*, and so on—are consistent with a multitude of alternative non-additive representations. Consider multiplicative measures, which map $\langle \mathbf{L}, \succsim; \circ \rangle$ not into $\langle \mathbb{R}^{\geq 0}, \geq; + \rangle$ but into $\langle \mathbb{R}^{\geq 1}, \geq; \times \rangle$ instead (see Hölder 1901; Krantz et al. 1971, 11–12, 99ff; Narens 1985, 27–31). Let the multiplicative (base 2) version of the meter scale be called the *schmeter scale*; it corresponds to a homomorphism m' that maps $\langle \mathbf{L}, \succsim; \circ \rangle$ onto $\langle \mathbb{R}^{\geq 1}, \geq; \times \rangle$ such that

a) $m'(L) \geq 1$ and $m'(L_m) = 2$
b) $L \geq L'$ iff $m'(L) \geq m'(L')$
c) $L \circ L' = L''$ iff $m'(L) \times m'(L') = m'(L'')$

On the schemeter scale, the value assigned to $L \circ L$ will always be equal to the square of the value assigned to L. Since 1 meter is 2 schmeters, then, 2 meters is 4 schmeters and 4 meters is 16 schmeters. Hence, if 4 meters is *twice as long* as 2 meters, it follows that 16 schmeters is *twice as long* as 4 schmeters.

The point is not that there's some contradiction here—there isn't. Rather, it's that the qualitative relation between $L \circ L$ and L is no more a *twice*-ish relation than it is a *square*-ish relation. In fact for *any* function $m^* : \mathbf{L} \mapsto \mathbb{R}$ that's related to the meter scale by an order-preserving transformation, there exists a numerical operation $*$ on $\mathbb{R}$ such that m^* is a homomorphism from $\langle \mathbf{L}, \succsim; \circ \rangle$ into $\langle \mathbb{R}, \geq; * \rangle$. Our use of 'twice as long' to refer to and describe the relation between $L \circ L$ and L reflects only a conventional preference for additive over alternative measures that are, from a mathematical point of view, all equally adequate to the task of measuring length.

But the conventionality runs deeper still, for it arises also in the *choice of qualitative system*. Again, length supplies a useful example. Earlier I characterised $\circ$ on $\mathbf{L}$ in terms of laying objects end-to-end. However, there are other concatenation operations involving lengths that we might have used as primitives instead. One alternative (also discussed by Ellis 1968, 80–1) is *right-angled concatenation*. Say that $L \odot L' = L''$ just when L'' is the length of the hypotenuse of the right-angled triangle with catheti of lengths L and L'. (See Figure 2.3.) *Right-angled concatenation* has all the same key properties as *end-to-end concatenation* that permit additive measurement. In an alternative history, then, we might have chosen to measure length by mapping $\langle \mathbf{L}, \succsim; \odot \rangle$ into $\langle \mathbb{R}^{\geq 0}, \geq; + \rangle$ (or $\langle \mathbb{R}^{\geq 1}, \geq; \times \rangle$, or...). This would have simplified how we express the relationships between sides of a right-angled triangle, though it would also have made calculating end-to-end concatenations of distances more difficult.

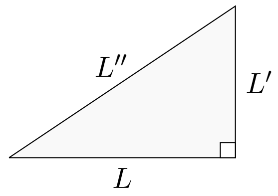

Figure 2.3: L'' is the right-angled concatenation of L and L' ($L \odot L' = L''$).

With all that said, I don't want to give the impression that anything goes. A minimal constraint on the choice of qualitative system is that the primitives must be *natural*. Without such a constraint, we trivialise the whole endeavour. For instance, assuming no more than that $\mathbf{L}$ can be mapped into $\mathbb{R}$, we know already that there exists a binary relation R on $\mathbf{L}$ that maps into $\geq$, in the sense that

$$(L, L') \in R \text{ iff } \varphi(L) \geq \varphi(L')$$

Likewise, supposing only that $\langle \mathbf{L}, \succsim \rangle$ maps into $\langle \mathbb{R}, \geq \rangle$, we know already that there will exist at least one ternary relation R such that

$$(L, L', L'') \in R \text{ iff } \varphi(L) + \varphi(L') = \varphi(L'')$$

So the fact that we can then find *some* relations on $\mathbf{L}$ corresponding to $\geq$ and $+$ is thoroughly uninteresting. We can trivially derive such relations from *any* mapping of $\mathbf{L}$ into $\mathbb{R}$, so long as we're permissive enough about what counts as a relation. It's a matter of convention what we take the primitive relations in our qualitative systems to be, but measurement is only *interesting* when those relations are *natural*.

2.5 Meaningfulness

A focus of the discussion to follow involves differentiating what is from what isn't meaningful in the measurement of belief. The most common strategy for drawing such distinctions uses *invariance* as a guide. Essentially, the idea is that a numerical property is meaningful only if it's invariant across alternative representations; otherwise it's a mere artefact of convention.

Compare the case of temperature. When represented in °C, water freezes at 0 and boils at 100, the hottest temperature recorded in Australia is almost exactly half way between these values (50.7), and more than double the hottest temperature recorded in Antarctica (19.8). But not all the numerical properties and relations just mentioned are *meaningful*. Measured in °F, water freezes at 32 and boils at 212, and the hottest recorded temperature in Australia (123) is less than twice the hottest temperature in Antarctica (68), though it'll still be just over half way between the freezing and boiling points of water. °C and °F are equally legitimate interval-scale representations of temperature—they're numerically distinct but they're not *meaningfully* distinct. The particular numerical values associated with each temperature and the ratios between those values vary between alternative scales, and are not, therefore, meaningful.

So far so good. But consider again the additive measures of $\langle \mathbf{L}, \succsim; \circ \rangle$. If φ and ψ are any two additive measures of that system, then

$$\varphi(L) = 2\varphi(L') \text{ iff } \psi(L) = 2\psi(L')$$

But this, too, is an artefact of convention. As we've just seen, the *qualitative* relation that holds between L and L' whenever L' is *twice as long* as L isn't *itself* a ratio relation in any deep sense, and if θ is a multiplicative measure then

$$\varphi(L) = 2\varphi(L') \text{ iff } \theta(L) = \theta(L')^2$$

In that broader sense, almost all the information in any numerical representation of $\langle \mathbf{L}, \succsim; \circ \rangle$ is an artefact of convention. There's approximately nothing that's invariant across *all* numerical representations of *any* qualitative system, and what *is* preserved is far too little to be of much interest.

The upshot is that meaningfulness needs to be understood relative to a fixed choice of numerical system. A more precise account of meaningfulness, and one that incorporates this lesson, originates with Pfanzagl (1968). I present it here in lightly modified form:

Definition 2.5 Suppose that $\mathbf{\Phi}(\mathcal{Q}, \mathcal{N})$ is non-empty, where $\mathcal{Q} = \langle \mathbf{X}, R_i; \bullet_j \rangle$ and $\mathcal{N} = \langle \mathbf{Y}, S_i; *_j \rangle$. For any $\varphi \in \mathbf{\Phi}(\mathcal{Q}, \mathcal{N})$ and any n-ary relation S on $\mathbf{Y}$, $R(S, \varphi)$ is the *relation induced on* $\mathbf{X}$ *by* S *and* φ if and only if

$$(x_1, ..., x_n) \in R(S, \varphi) \text{ iff } (\varphi(x_1), ..., \varphi(x_n)) \in S$$

S is $\mathcal{Q}$*-meaningful relative to* $\mathcal{N}$ when $R(S, \varphi)$ doesn't depend on the choice of φ in $\mathbf{\Phi}(\mathcal{Q}, \mathcal{N})$.

Where the intended $\mathcal{Q}$ and $\mathcal{N}$ are obvious given context, we simply say 'meaningful'. Note that if S is among the primitive relations S_i of $\mathcal{N}$, then $R(S, \varphi)$ is just the corresponding primitive relation R_i in $\mathcal{Q}$ and thus automatically $\mathcal{Q}$-meaningful relative to $\mathcal{N}$. So we're only interested in the case where S *isn't* among the primitives relations of $\mathcal{N}$. Note also that if S corresponds to one of the primitive *operations* in $\mathcal{N}$, then it *doesn't* automatically follow that S is meaningful, unless every homomorphism in $\mathbf{\Phi}(\mathcal{Q}, \mathcal{N})$ is a strong homomorphism.

It's helpful to compare cases where a numerical relation isn't meaningful. Observe first of all that, where $\mathbf{\Phi}(\mathcal{Q}, \mathcal{N})$ is non-empty, every numerical relation S will induce a corresponding relation $R(S, \varphi)$ on the qualitative system, relative to φ. So, when any ordinal scale φ maps $\mathcal{Q} = \langle \mathbf{L}, \succsim \rangle$ into $\mathcal{N} = \langle \mathbb{R}, \geq \rangle$, the 2:1 ratio relation induces a corresponding relation on $\mathbf{L}$ that holds for L, L' whenever $\varphi(L) = 2\varphi(L')$. But that relation isn't $\mathcal{Q}$-meaningful relative to $\mathcal{N}$, precisely because $R(2{:}1, \varphi)$ needn't equal $R(2{:}1, \psi)$ for every other ψ in $\mathbf{\Phi}(\mathcal{Q}, \mathcal{N})$. By contrast, the 2:1 ratio *is* meaningful relative to the additive measures of $\langle \mathbf{L}, \succsim; \circ \rangle$, since $R(2{:}1, \varphi)$ equals $R(2{:}1, \psi)$ for any two additive measures φ and ψ. Why does that matter? Because if the 2:1 ratio is meaningful with respect to the additive measurement of length, then we can draw generalisations and formulate laws involving that relation without worrying that it all depends on an arbitrary choice of scale.

In general, the idea is that a numerical relation is meaningful inasmuch as it always corresponds to the same qualitative relation regardless of what homomorphism we care to use, given a fixed choice of numerical system. That's just what 'meaningful' *means* in this context: *always picks out the same thing independent of the choice of scale.* So to make some headway on the matter of what should be considered meaningful in our numerical representations of belief, we need to say more about the kinds of qualitative structures that these representations are supposed to be representations of.

Chapter 3

Clarifications and Desiderata

The central questions to be addressed by an account of the measurement of belief are, in relation to a given purported numerical representation of belief: what is the qualitative system being represented, what is the numerical system in which it's represented, and under what conditions are such representations possible?

An epistemic approach to answering these questions, I said, is one that appeals to the internal structure of the belief system. Better: an epistemic approach is one according to which the qualitative system being represented can be characterised fully in terms of doxastic states and the relations between them, where a *doxastic state* is any type of mental state that has a belief-*ish* flavour. This might include states of all-or-nothing belief, particular levels of absolute confidence, judgements of relative probability, judgements of when one thing is evidence for another thing or when they are evidentially independent, holistic belief systems, and so on. In sum they are the sorts of things that have a mind-to-world direction of fit, broadly construed; or the sorts of things that reflect our opinions regarding what the world is like and what it might be like, and that ought to be directly responsive to evidence independent of our preferences. Epistemic approaches are covered in Chapter 4 and Chapter 5.

Decision-theoretic approaches instead appeal to relations between doxastic states and conative states (states with a desire-*ish* flavour) to explain what the numbers mean. Roughly, a paradigm decision-theoretic approach is one where the qualitative system is comprised of (formalised representations of) a conjoint system of beliefs and desires, related via their joint determination of a preference relation over a space of actions according to some decision rule; the numerical representation is then constructed to capture the relations holding between the three. These are covered in Chapter 6.

Before we delve into the details, though, this chapter provides some background clarifications on what a theory of belief measurement is and what it's not (§3.1 and §3.2), followed by a some assumptions (§3.3) and general desiderata (§3.4) that will be relevant to the discussion throughout.

3.1 Quantitation, not elicitation

In the classic presentations of the RTM, qualitative systems are usually understood to be *empirical relational systems*. These are systems built around primitives that are directly and publicly observable in the context of some experimental procedure. For example, rather than characterising the length system $\langle \mathbf{L}, \succsim; \circ \rangle$ as a set of attributes and higher-order relations between them, if I were doing things in the classical manner then I'd have characterised it instead as a set of physical objects, the observable *at least as long* relation between them, and a concatenation operation interpreted as the physical process of taking two objects and joining them to form a new composite object. Essentially, empirical relational systems are systems in which nothing is hidden from view—the relations should be open to observation, the relata should be things we can touch and see and poke and prod, and the operations are procedures on or processes observed in the entities being measured.

There are some problems that arise when measurement is understood this way. (These problems were not unknown to the founders of the representational theory; cf. Krantz *et al.* 1971, 27–31.) For example, in violation of transitivity axioms it will be possible to have a series of objects, each longer than the preceding by an imperceptible amount, such that adjacent objects will be observed to be of the same length even while the last is much longer than the first. Similar problems arise for most empirical relational systems, and all point to the same basic issue: quantities cannot be characterised in terms of the experimental procedures by which they're measured, for the simple reason that no such procedure is perfect. Instead, measurement procedures are developed on the basis of what our theories imply about the conditions under which observable experimental outcomes will reliably (albeit imperfectly) correlate with variations in some limited range of magnitudes of the quantity we desire to measure.[1]

Such will be familiar from the history of operationalism. But it would be a mistake to dismiss classical RTM's focus on empirical relational systems as an offshoot of that now-defunct philosophy. More illuminating to say instead that the mathematical framework of the RTM was developed to play two distinct explanatory roles. On the one hand, it was to help explain how we can use numerical properties and relations to represent and reason about bits of the world that aren't themselves numerical in nature. This explanation appeals to structure-preserving mappings between qualitative and numerical systems, and *observability* is entirely irrelevant here. On the other hand, the same formalisms were intended to supply necessary conditions for the design of empirical measurement procedures. An empirical relational system for the measurement of mass, for example, was supposed to be formulated such that it could be directly implemented in some empirical measurement procedure; hence the pervasiveness of error in all actual measurement practices was taken to present a serious problem for the RTM (cf. Krantz *et al.* 1971, 1–9, 25, 27–31).

[1] See (Mari *et al.* 2017) for relevant discussion and a detailed account of the theory-based construction of one such procedure for the measurement of the mass of stars.

We can—and should—keep these roles separate. The RTM is great for the first, not so great for the second.[2] As Kyburg once said, the 'theory of measurement is difficult enough without bringing in the theory of making measurements' (1984, 7). Unfortunately, ambiguity in how we use the term 'measurement' can easily obscure this point. Compare the "measurement of mass" *qua* abstract pairing of determinate mass attributes with numbers such that relations between the latter usefully mirror relations between the former; versus the "measurement of mass" *qua* empirical procedure for determining the mass of particular objects by means of, say, an equal-armed pan balance. The original intention for the RTM was to be a theory of both. The unfortunate result has been that it's routinely criticised for being of little relevance to the actual measurement practices of working scientists (e.g., Borsboom 2005; Mari 2005; Reiss 2016), when it's much better understood as a framework for understanding meaning and meaningfulness in our numerical representations of systems of determinable attributes as posited by a scientific theory.

In light of this, let me emphasise that a theory of belief measurement *as presently understood* is not in the business of explaining how we might gather empirical evidence as to the strength of an agent's beliefs through the observation of their behaviour, nor how we might elicit their beliefs by any other means. Mario Bunge (1973) once recommended avoiding the ambiguity of 'measurement' by referring to the abstract sense as *quantitation*. In those terms, the topic here is the quantitation of beliefs, not their elicitation.

Consequently, we ought make no assumptions about the observability of the qualitative primitives posited within a theory of belief measurement. These systems posit psychological relations—things like *is more confident than*, *is more desirable than*, *is indifferent between*—and I think it would be an error to assume that such relations will generally be directly observable in behaviour.[3] It would be a deeper error still to assume that these relations *must* to observable if we're to justify theses about the structure of the qualitative systems involving them. Quantities are posits of our scientific theories, and like any other posits they need not be directly observable. The justification for the thesis that a particular qualitative system has a certain formal structure that permits numerical representation in such-and-such a manner needn't derive from any direct observation of that structure, but can instead derive indirectly from the empirical and theoretical virtues of the theories that presuppose a system of quantities endowed with that structure. In this respect the measurement of belief is no different than the measurement of any other quantity.

[2] This opinion neither new nor uncommon; similar can be found expressed in Roberts (1985), Mundy (1987; 1994), Swoyer (1991), Narens & Luce (1993), Decoene *et al.* (1995), Mari *et al.* (2017), and Bacelli (2020). See also (Michell 2021) for a useful overview of the history of thought on this matter.

[3] Perhaps, under special circumstances, some limited part of a person's total preference ordering might be directly 'revealed' through their choice behaviour and choice dispositions. Many people have thought so. I doubt it. At best there's a defeasible evidential relationship between choice and preference, and the connection is too loose to say that the latter are ever directly observable via the former.

3.2 Measurement, not metaphysics

This is an essay on measurement, not metaphysics. The topic is not quite *what are degrees of belief?*, but *what do our purported numerical representations of belief actually represent and how do they represent it?* These can be hard to keep separate, but separate them we should—lest we end up rejecting entirely reasonable accounts of the measurement of belief by conflating them with hideously implausible accounts of the metaphysics of belief. This is a problem especially for decision-theoretic approaches, due in large part to a historical association with behaviourist (and behaviourist-lite) metaphysical theories that purport to reduce beliefs to little more than preferences as revealed by choice.

A metaphysics of belief, I take it, is concerned primarily with the kinds of ontological and/or conceptual dependence that hold between doxastic states of different kinds, and between doxastic states and the wider world.[4] The core task of such a metaphysics is, in short, to explain what kinds of doxastic state-types exist and where they ought to be situated relative to one another and relative to the rest of the world within some general conceptual framework and/or global ontology of the universe.

One major metaphysical division is between *realist* and *anti-realist* views; the former broadly speaking being the idea that the correct attribution of a doxastic state to an agent depends on objective facts about those agents, and the latter saying instead that correct attribution depends somehow on who's doing the attributing. Some versions of *interpretivism* or *instrumentalism* fall into the anti-realist camp; such will typically say that an agent's beliefs are just those an interpreter can usefully employ to explain the agent's behaviour, or perhaps her behavioural dispositions conditional on evidence. Among realists, another major division is between *representational* and *non-representational* theories. The former explains what it is to have a doxastic state with such-and-such content by connecting those states to the presence of some internal mental representation of that content. Non-representational theories link doxastic states instead to not-necessarily-representational states of the agent that are systematically related to the contents thereof. Among non-representational views are *behaviourist* theories, which analyse doxastic states as certain patterns of behaviour; *dispositionalist* theories, which tie doxastic states to a suite of associated (and not-necessarily-behavioural) dispositions; and *functionalist* theories, which propose to analyse doxastic states by reference to a functional role that typically revolves around relations between beliefs over time given evidence and between beliefs and desires in the production of behaviour.

[4] It's not easy to say precisely what ontological and conceptual dependence relations are, and any characterisation I give will be subject to debate. At least roughly, a concept C is *conceptually* more fundamental than another concept C' when C' can be analysed in terms of C but not vice versa; and a property (or state-type) P is *ontologically* more fundamental than another P' when the instantiation of P' necessarily depends on the instantiation of P but not vice versa. Another way to distinguish the two is via their distinct explanatory roles: ontological dependence helps explain necessary connections between properties, while conceptual dependence relations explain a priori connections between concepts.

The core questions dealt with by a theory of belief measurement, by contrast, concern the narrower issue of quantitation. In relation to a purported numerical representation of some doxastic state or set of such states,

1. what is the qualitative system $\mathcal{Q}$ being represented,
2. what is the numerical system $\mathcal{N}$ in which it's represented, and
3. under what conditions are such representations possible?

Such questions aren't orthogonal to the metaphysics of belief, but they're not really aimed at the same sorts of issues either. Consequently, each and every one of the approaches to the measurement of belief that I'll discuss in the chapters that follow is compatible with a very wide range of perspectives on the metaphysics of belief, including any and all of the perspectives just mentioned. There are some ties between the two topics, and some theories of measurement will naturally fit better with some metaphysics and vice versa, but in general one cannot read metaphysics off of measurement.

That should be clear enough in the case of epistemic approaches. But experience teaches that it can be very easy to slip into a 'metaphysical' interpretation of the representation theorems that form the foundations for decision-theoretic approaches, and so I want to consider that case in a bit more detail. I'll start with an oversimplified sketch of a decision-theoretic approach involving the conjoint measurement of beliefs, desires, and preferences.

Let $\mathcal{G}$ be a set of simple binary gambles of the form "receive c if p is true, and c' otherwise", and let $\succsim$ be a preference relation over $\mathcal{G}$. We suppose that the agent's preferences over these gambles are determined by the expected utility rule relative to her degrees of belief and the strengths of her desires for the consequences. If we let $\mathcal{A}$ be a set of propositions, and $\mathcal{C}$ be a set of consequences, then each gamble can be modelled as the 3-tuple (c, p, c') in $\mathcal{C}\times\mathcal{A}\times\mathcal{C}$. Thus, we are searching for a three-vector homomorphism φ from $\langle\mathcal{C}\times\mathcal{A}\times\mathcal{C}, \succsim\rangle$ into $\langle\mathbb{R}\times\mathbb{R}\times\mathbb{R}, \geq\rangle$ that 'decomposes' into a representation of beliefs $\beta : \mathcal{A} \mapsto [0, 1]$ and a representation of desirabilities for consequences $\delta : \mathcal{C} \mapsto \mathbb{R}$ such that $(c_1, p, c_2) \succsim (c_3, q, c_4)$ iff

$$\delta(c_1)\beta(p) + \delta(c_2)\big[1 - \beta(p)\big] \;\geq\; \delta(c_3)\beta(q) + \delta(c_4)\big[1 - \beta(q)\big]$$

Suppose that $\succsim$ satisfies the conditions required for the desired representation to exist. The result is a conjoint representation of beliefs, desires, and preferences, constructed in the first instance to capture the relationship between beliefs and desires in connection to preferences.

This kind of representation is associated with the idea that beliefs and desires are nothing over and above preferences—that *what it is* to have such-and-such beliefs and desires *just is* to have such-and-such preferences. On reflection, though, that association is quite unusual. If I describe a structure for the conjoint measurement of momentum as determined by mass and velocity, no one leaps to the conclusion that mass and velocity are ontologically dependent on momentum. Likewise, if I describe a conjoint measurement structure for overall value as determined by monetary value and sentimental value, no one infers that our concepts of *monetary value* and *sentimental value* ought to be analysed in terms of our concept of *overall value*. Such inferences would be fallacious, and quite obviously so. Yet the parallel inferences are frequently made in connection to the conjoint measurement of beliefs, desires, and preferences.

According to the decision-theoretic approach, what the conjoint representation teaches us is that some lawlike relation or relations between beliefs and preferences is *explanatorily relevant to the quantitation of belief*. Such in no way implies that preferences are conceptually or ontologically more fundamental than beliefs. The explanatorily relevant relations between belief and preference may not be dependence relations at all. For example, the kind of decision-theoretic approach I outlined would be consistent with a functionalist metaphysics according to which beliefs, basic desires and preferences are interrelated posits in a psychological theory, such that none are reducible to the others, and such that their statistically or biologically normal causal interactions can be systematically represented within an expected utility framework.

Observe, also, that such a functionalist might say that the relation between beliefs and preferences is critical for explaining the quantitation of belief, even while saying that the characteristic role of belief isn't *exhausted* by that relation. One might reasonably suppose that an important part of the functional role of belief concerns the relationship between beliefs and sensory evidence—a state cannot rightly be said to "play the belief-role" if it isn't appropriately sensitive to perceived changes in the environment. Such connections will be crucial when providing a functionalist *analysis* of what beliefs are, but that doesn't imply that they need also be mentioned in an explanation of why it makes sense to represent a system of beliefs within a certain numerical framework.

(Compare the case of mass. Our concept of *mass* can be plausibly analysed in terms of its role within contemporary physics: mass is the property that does the best job of satisfying the total theoretical role associated with 'mass'. But mass does many things. The mass of an object is proportionate to its resistance to acceleration as measured by an observer at rest with respect to it. It's also proportionate to the strength of the gravitational field the object exerts on others, and its total rest energy. Mass is tied to momentum and velocity, density and volume, and to how fast a transverse wave travels through a string attached to a fixed point at each end. Mass also plays a role in our theories of stellar evolution; for instance, a white dwarf with mass exceeding about 1.4 solar masses will succumb to electron degeneracy pressure and collapse into either a neutron star or a black hole. So if you want to analyse *mass* in terms of its total theoretical role, there's a lot you need to mention. But if you want to give an explanation of why it makes sense to measure mass on a ratio scale, then not all of that is going to be necessary or relevant. The relations we use to analyse the concept of a quantity can come apart from the narrower class of relations we use to explain the quantitation of that quantity.)

An account of the measurement of belief just isn't in the business of explaining ontological or conceptual dependence relations that hold between different kinds of doxastic states, nor between doxastic states and non-doxastic states. It would be wise, then, to be very careful when drawing metaphysical conclusions from measurement-theoretic premises.

3.3 Simplifying assumptions

Having said some things about the sorts of things we shouldn't be assuming, let me talk about the assumptions I will be making. The first two are simplifying assumptions about how we model contents:

Assumption 1. Degrees of belief have propositional contents, where propositions can be modelled as subsets of some non-empty set of possible worlds (henceforth denoted Ω).

Assumption 2. For each agent and all propositions p, there exists an algebra of propositions $\mathcal{A}$ on Ω such that the agent has some degree of belief towards p iff p belongs to $\mathcal{A}$.

By 'possible', I mean at least consistent with classical logic. An algebra of propositions is defined like so:

Definition 3.1 $\mathcal{A}$ is an *algebra of propositions* on Ω iff it is a non-empty set of subsets of Ω, and for all $p, q \subseteq \Omega$,

1. If p is in $\mathcal{A}$, then $\Omega \setminus p$ (henceforth $\neg p$) is in $\mathcal{A}$
2. If p and q are in $\mathcal{A}$, then $p \cup q$ is in $\mathcal{A}$

Furthermore, an element $a \in \mathcal{A}$ is an *atom* of the algebra iff $a \neq \varnothing$ and for every $p \in \mathcal{A}$, either $a \cap p = a$ or $a \cap p = \varnothing$.

These are substantive assumptions indeed, and I'm not so confident they're true—but they're also both very standard assumptions in the present context and each does a great deal to help simplify a great many matters. Also for the sake of simplicity, I'll mostly focus on measurement structures involving finite algebras. This is not because I think that agents can have degrees of belief towards only finitely many propositions, but because trying to cover infinite cases would add significant complexity with little by way of philosophical pay-off.

I should say a bit more about these assumptions, though, since they'll play an important role at some points (especially in connection to logical omniscience, §4.2). A consequence of **Assumption 1** is that the contents of belief are *coarse-grained*: if p and q are logically equivalent, then $p = q$. More generally, where $p \Rightarrow q$ means that p logically implies q, then it follows from **Assumption 1** that $p \subseteq q$ iff $p \Rightarrow q$. But I did not refer to the assumption as 'simplifying' because of this—there's a lot to be said in favour of coarse-grained content! (e.g., Stalnaker 1984; Lewis 1986; Chalmers 2011.) Rather, it simplifies because it ignores *de se* content and common strategies for the representation thereof that require going beyond possible worlds (e.g., Lewis 1979).

Opponents of coarse-grained content often suppose that we can model more fine-grained belief contents using sets of possible *and* impossible worlds. Roughly, the idea is that wherever we want to differentiate between two logically equivalent contents p and q, we include an impossible world where one of these holds but the other doesn't; hence the set of p-worlds will come apart from the set of q-worlds. More generally, in this context the $\subseteq$ relation need no longer track the $\Rightarrow$ relation, which aids in dealing with the many problems of logical omniscience.

But matters are not quite so easy. One cannot simply throw a bunch of impossible worlds into Ω without breaking something elsewhere, especially in the context of **Assumption 2**. To explain why, I'll need to say a bit more about what impossible worlds are and how they're used to model contents. I'll adopt the modal ersatz approach found in (Nolan 1997), though essentially the same points can be made for other popular accounts of impossible worlds (e.g., linguistic ersatzism or extended modal realism, see Elliott 2019b for discussion).

We take *propositions*—the potential objects of our beliefs and the meanings of our declarative sentences, whatever they may be—to be ontologically primitives. Given that, we let a world ω be any set of propositions, and we say that p is true at a world ω just in case $p \in \omega$. There is, of course, a one-to-one correspondence between each primitive proposition p and the set of worlds containing p (the p-worlds). We say that a world is possible just in case it's *complete* (contains either p or its negation, for every proposition p) and *consistent* (has no logically inconsistent subsets); otherwise, it's impossible. If Ω contains only possible worlds, then $p \Rightarrow q$ implies that every p-world in Ω will be a q-world. But if Ω isn't restricted to possible worlds, then it may be that $p \Rightarrow q$ even while there are some impossible p-worlds in Ω that aren't q-worlds. Much therefore depends on what kinds of worlds get to go into Ω; the richer the space of worlds, the more distinctions we can draw between logically-equivalent contents modelled as sets of worlds and the looser the link between $\Rightarrow$ and $\subseteq$.

Impossible worlds theorists will often assume a very rich space of worlds characterised by an *unrestricted comprehension* principle: for any complete set of propositions $\mathcal{P} = \{p, q, \ldots\}$, there is a world $\omega \in \Omega$ such that $\omega = \mathcal{P}$. Roughly, for any possibility or impossibility, there's a world that verifies it; and the principle ensures that some p-worlds aren't q-worlds even where $p \Rightarrow q$. However, it also has the consequence that a great many subsets of Ω are *meaningless*. These are sets of worlds that correspond to no proposition whatever; there is nothing that's true at all and only the worlds in a meaningless set, they are instead just artefacts of the construction of contents are sets of sets of propositions. As Nolan (1997, 563) points out, given *unrestricted comprehension* any set of possible worlds will be meaningless in this sense. For any set of possible worlds $\{\omega_1, \omega_2, \ldots\}$ there will be some propositions they all have in common. Given that, let ω_i be a world such that everything true at all the worlds in $\{\omega_1, \omega_2, \ldots\}$ is also true at ω_i, but the negation of one or more of those things is also true at ω_i. It follows that ω_i is an impossible world. So, there is nothing true at all *and only* a set of possible worlds—such sets are meaningless.

The existence of meaningless subsets of Ω isn't intrinsically problematic. However, it does not play nicely with **Assumption 2**. An algebra of propositions is closed under relative complements and binary unions, and in the presence of *unrestricted comprehension* two facts follow. First, the relative complement of any meaningful proposition is meaningless: for any p and q there will be worlds where both are true, hence there's no q such that the set of p-worlds doesn't intersect with the set of q-worlds. Second, the union of any two meaningful propositions is meaningless: for any p and q there can be no r such that the set of r-worlds is the union of the p-worlds and the q-worlds, since then every

p-world would be an r-world but for any p and r there will be some p-worlds that aren't r-worlds. In short, then: any algebra of propositions defined on a sufficiently rich space of possible and impossible worlds will consist *mostly* of meaningless sets—and we shouldn't want to represent agents as having beliefs towards entities that correspond to no proper object of belief.

You might think there's an easy response: the main premises are the *unrestricted comprehension* principle and **Assumption 2**, so we can simply deny one or both of those premises. Again, though, matters aren't so simple. The *unrestricted comprehension* principle is required for the most attractive results that impossible worlds are advertised have in relation to fine-grained content and logical omniscience (see Nolan 2013 for an overview). Moreover, it's a mistake to suppose that *unrestricted comprehension* is necessary for the conclusion—as if the issue would simply disappear were we to adopt a more restricted principle. As shown in (Elliott 2019b), the real problem is that **Assumption 2** imposes a Boolean algebraic structure over meaningful subsets of Ω, which forces the worlds in Ω to conform to a Boolean logic. Under quite minimal richness conditions on what kinds of *possible* worlds go into Ω, either (a) every algebra of sets on Ω will contain meaningless sets of worlds or (b) the worlds in Ω will be closed under the $\{\neg, \wedge\}$-fragment of Boolean logic (or something to the same effect).

Nor is it easy to deny **Assumption 2** since—as we'll see—many approaches to the measurement of belief make important use of that assumption. This includes all of the epistemic approaches we discuss, and a number of decision-theoretic approaches. The reason boils down to the fact that representation of any quantity on more than a mere ordinal scale requires a qualitative structure richer than what can be provided by a single weak ordering over the magnitudes thereof—some additional relation will be required for the extra-ordinal structure of the numerical representation to grock on to. Thus, for example, in the measurement of length we require not only the *at least as long* relation, but also a concatenation operation that can be mapped into addition. Likewise for conjoint measurement, where the additional structure is supplied by reconstructing $\succsim$ on the single quantity $\mathbf{C}$ as a quarternary relation over $\mathbf{A}\times\mathbf{B}$ and then using induced relations between the factors $\mathbf{A}$ and $\mathbf{B}$ to supply the additional structure for the representation. For theories of belief measurement—and especially for epistemic approaches—the additional structure that allows for more-than-merely-ordinal measurement is typically characterised by set-theoretic relations between contents, and in such a way that presupposes the algebraic structure guaranteed by **Assumption 2**.

The point here is *not* that there's no hope for impossible worlds. The point, rather, is that incorporating impossible worlds into contemporary theories of belief measurement will at least require careful consideration about the nature of content and likely some adjustments to the formal models and associated proofs. The common thought is that impossible worlds present an *easy* solution to the problems of coarse-grained content and logical omniscience—just throw some impossible worlds into the mix and you're done. It is not so easy. In that sense, then, the conjunction of **Assumption 1** and **Assumption 2** can be considered a simplifying assumption, too.

One more simplifying assumption:

Assumption 3. Degrees of belief are precise.

I don't think this is realistic. Imagine, for instance, that a down-trodden magician has just rolled into town. He has a coin, which you happen to know is biased but you know not in what direction the bias lies nor to what degree; and an old deck of cards with some unspecified number of cards missing. The magician tosses the coin and pulls out a single card from the deck. For p and q as follows, do you have at least as much confidence in p as in q, at least as much in q as in p, or exactly the same degree of confidence in both?

$$p = \text{the coin will land heads} \qquad q = \text{the card will be red}$$

If degrees of belief are precise (represented by real values), then there must be a fact of the matter. Plausibly, there isn't, or at least needn't be.

Over the past few decades especially, something of a consensus has emerged regarding the representation of 'imprecise' degrees of belief (see, e.g., Walley 1991; Kaplan 2010; Joyce 2010). Instead of the real-valued functions employed in classical models of graded belief, we instead use a set of real-valued functions—a *credal set*, or as it's often known in philosophy, a *representor*. The rough idea is that what a representor represents is what's true according to all functions in the set. Thus, for example, we say that the subject has at least as much confidence in p as she does in q just when every function in her representor assigns a value to p that's at least as great as the value assigned to q.

A general strategy exists for the construction of these 'representor' representations that can be applied to the different the epistemic and decision-theoretic approaches discussed below (e.g., Evren & Ok 2011; Alon & Lehrer 2014; Alon & Schmeidler 2014; Augustin *et al.* 2014; Hawthorne 2016). The main move is to replace the common *weak order* axiom they all employ with a weaker *preorder* axiom, thus allowing for incompleteness in the primitive psychological relations being represented. The 'imprecise' representation is then constructed from the precise representations of the possible completions of that preorder. But since this *is* a general strategy that works more or less the same way across epistemic and decision-theoretic approaches, I've neglected to include details. Instead, I'll take it as read that the real-valued measures of belief considered below are idealisations. (I'll have a little more to say about this in §6.4.)

3.4 Desiderata

The remainder of this chapter will outline four basic desiderata for a theory of belief measurement. Again I'll need to start with some terminology.

According to the thesis of *probabilism*, ideally rational agents are those whose beliefs can be accurately represented by some probability measure. Probability measures are typically defined as follows:

Definition 3.2 Where $\mathcal{A}$ is an algebra of propositions on Ω, $\mu : \mathcal{A} \mapsto \mathbb{R}$ is a *probability measure* iff, for all $p, q \in \mathcal{A}$,

1. $\mu(p) \geq 0$ (*non-negativity*)
2. $\mu(\Omega) = 1$ (*normalisation*)
3. If $p \cap q = \varnothing$, then $\mu(p \cup q) = \mu(p) + \mu(q)$ ($\sqcup$-*additivity*)

Exactly *what it is* for a system of beliefs to be represented by a probability measure is, of course, a question to be settled by an account of the measurement of belief. But set that aside for the time being. A weaker version of the thesis, what Kaplan (2010) calls *modest probabilism*, requires only that an ideally rational system of beliefs can be represented by a non-empty set of probability measures. Finally, according to what I'll call *really modest probabilism*, at least some rational systems of belief are represented by probability measures.

While there are occasional arguments against probabilism, these typically call for exceptions to the idea that ideally rational agents *must* be represented by (sets of) probability measures. So I take it that really modest probabilism ought to be innocuous, and as such we should want a theory of belief measurement that's able to make sense of it:

Desideratum 1. A theory of the measurement of belief should be consistent really modest probabilism.

That is, the theory should explain how it's possible for a system of beliefs to be accurately represented by a probability measure (or a set thereof). Since this is what the large majority of epistemic and decision-theoretic approaches already in fact do, I don't expect much resistance on this front. Still, especially when combined with the remaining desiderata it plays an important role in constraining what counts as a desirable theory of belief measurement.

For the next, let's say that a measure of belief is *cardinal* (as opposed to *merely ordinal*) if it's unique up to something stronger than an order-preserving transformation. So, for example, interval-scale and ratio-scale measures will count as cardinal measures in this sense. Given that,

Desideratum 2. Probabilistic representations of belief are (at least in some theoretical contexts) cardinal measures of those beliefs.

One very simple reason to accept this is intuition. Most will be happy to say that a rational agent ought to have about 50% confidence that a fair coin will land heads on a single toss, which should be *half as much* confidence as they have regarding it landing either heads or tails, and *twice as much* confidence as they ought to have regarding it landing heads twice in a row. Or, more simply still, it's clearly sensible to say that a person can have *much more* confidence in one thing than in another. Such makes sense only if beliefs are measurable on something stronger than an ordinal scale.

I'm inclined to take these intuitions seriously, as strongly indicative of how we pre-theoretically think about confidence. But I wouldn't want to rest my case on intuition alone. A stronger reason to accept **Desideratum 2** arises from the fact that cardinal information has a theoretical role to play in our theories of rational decision-making. Consider the following example. We imagine first that Ramsey has to choose between two gambles:

α: receive \$1 if p is true, nothing otherwise
β: receive \$2 if p is false, nothing otherwise

Suppose also that Ramsey considers p less probable than Ω but more probable than $\neg p$. Without loss of generality, let $\mathcal{A} = \{\Omega, p, \neg p, \varnothing\}$. A probability measure will be an ordinal representation of Ramsey's confidences just in case it assigns a value to p that's strictly between 1 and $\frac{1}{2}$. As such, there will be at least two ordinally-equivalent probability measures, μ_1 and μ_2, such that

$$1 > \mu_1(p) > \frac{2}{3}, \qquad \frac{2}{3} > \mu_2 > \frac{1}{2}$$

If confidence is measured on nothing stronger than an ordinal scale, there should be no difference in meaning between μ_1 and μ_2. However, there *is* a decision-theoretic difference between them: Ramsey should weakly prefer α iff his confidence in p is more than twice his confidence in $\neg p$. At that point, the higher probability of winning with α outweighs the promise of a larger prize with β. So expected utility theory is inconsistent with the thesis that confidence is measured on an ordinal scale (I'll say more about this in §6.3.)

It's easy to verify that the same holds for the vast majority of alternatives to expected utility theory, including normative theories (designed to represent the decisions of ideally rational agents) and descriptive theories (designed to represent the decisions of more realistic agents). And we needn't rest the case on decision-theoretic examples either—much the same holds in contemporary epistemology, where a great deal of theory and argument presumes the more-than-merely-ordinal measurement of belief. Three very brief examples; I'm sure if you search you will find more. First, the relation of probabilistic independence is important for Bayesian theories of evidence and learning, but independence relations can vary between ordinally-equivalent numerical representations (see §5.3). Second, epistemic utility theory appeals to numerical properties that differentiate ordinally-equivalent probability measures (see Mayo-Wilson & Wheeler 2019, p. 19). Third: there are meaningfully distinct but ordinally-equivalent objective chance functions, and as a result the Principal Principle (Lewis 1980) presupposes distinctions in rational belief that cannot be captured on an ordinal scale. In sum: if our numerical representations of belief are to play the theoretical roles they are generally taken to play in rational belief and rational decision-making, then they cannot be mere ordinal-scale measures.

Together, these two desiderata imply that at least some agents have beliefs representable by a probability measure, and that such measures aren't ordinal scales. The next is aimed instead at extending the same to non-ideal agents:

Desideratum 3. Logical omniscience is not a prerequisite for the cardinal measurement of belief.

In the present context, I take it that an agent is *logically omniscient* just in case, if p logically implies q, then the agent has at least as much confidence in q as they do in p. Logical omniscience implies, among other things, (i) the agent never has more confidence regarding any logically impossible proposition than they do for any other proposition, (ii) they never have more confidence in any proposition than they do in any logically necessary proposition, and (iii) if two propositions are logically equivalent, then the agent has the same degree of confidence in each. The third is, of course, also a consequence of **Assumption 1**.

The argument I'll provide for **Desideratum 3** *is* based on intuition—it's hard to believe that an agent who makes some logical mistakes is thereby automatically precluded from believing one proposition *much more* than another, or *about half as much* as another, and so on. I'm not ideally rational, and neither are you. We are less-than-ideally rational, and one likely manifestation of our lack of this is that we aren't logically omniscient. Given **Assumption 1**, any probability measure automatically determines a logically omniscient probability ordering, so it follows that our beliefs cannot be accurately represented by any probability measure (or set thereof) defined on an algebra of propositions on Ω. But this doesn't prevent us from believing propositions to different degrees, where those are measured on something stronger than a mere ordinal scale.

The final desideratum is an anti-disjunctiveness condition:

Desideratum 4. A theory of belief measurement should not be fundamentally different for ideal versus non-ideal agents.

If we're going to say that both ideally rational and non-ideally rational agents can have degrees of belief that are measured on something stronger than an ordinal scale, then we should also want an explanation that makes sense in both cases—a unifying theory is a better theory. There doesn't appear to be any difference in *meaning* when we say (e.g.) that Jules is *much more* confident that the world is round than that it's flat, depending on whether Jules is ideally rational or less-than-ideally rational. If that intuition is right, then fundamentally the same explanation of quantitatability should apply in either case.

I intend for this to be compatible with the idea that there might be more than one adequate approach to the measurement of belief. It might be, for example, that a decision-theoretic approach is apt for the purposes of decision theory, and that an epistemic approach is apt for other theoretical contexts, with no fact of the matter as to which is the correct way of doing things. It's not unusual that there might be complementary ways of explaining the quantitation of a given quantity. For example, the ratio-scale measurement of mass can be explained as an instance of fundamental extensive measurement, or conjoint measurement, or, given an appropriate choice of base quantities, derived measurement—there is no fact of the matter as to which is correct. But the key term there is *complementary*. The various ways to explain the quantitation of mass are not disjunctive in the sense of giving one explanation for how mass is measured that applies to certain cases, and a fundamentally distinct and incompatible explanation for other cases. That's the kind of disjunctiveness we should avoid.

A final note. The joint effect of these four desiderata will be that we should want to find a non-disjunctive theory of how beliefs can be measured on something stronger than an ordinal scale that's consistent with really modest probabilism, but at the same time isn't restricted to representing only the beliefs of agents who satisfy the over-demanding requirements of logical omniscience. But

I'll not be explicitly evaluating the theories of belief measurement considered below in terms of those desiderata. Evaluation is left primarily to the reader, and you may take issue with some or all of what I take to be desirable in a theory of belief measurement. Rather, the desiderata are offered by way of explanation for why I've chosen to focus on certain topics in the chapters that follow: the meaningfulness of extra-ordinal information, probabilistic and non-probabilistic representations, and logical omniscience.

Chapter 4

Epistemic Approaches: Comparative Confidence

The most straightforward and best-known epistemic approach involves the probabilistic representation of (complete) binary comparative confidence relations. For ease of reference, I'll refer to this as the *standard epistemic approach.* This chapter begins with an overview of the standard epistemic approach (§4.1), after which we consider the problem of logical omniscience and non-probabilistic generalisations (§4.2). Several alternative epistemic approaches are also considered in the next chapter.

For the entirety of the present chapter, let $\succsim$ be interpreted relative to some agent α, and we read $p \succsim q$ saying that α has at least as much confidence in p as she has in q. Supposing $\succsim$ is a weak order, it's then natural to interpret $\succ$ as a *more confidence* relation, and $\sim$ as an *equal confidence* relation. Where $p \sim q$, I'll sometimes say that p and q are *equiprobable*; this usage shouldn't be understood to presuppose that $\succsim$ has a probabilistic representation.

4.1 Probabilistic representations

The main results in this area concern the conditions under which a system comprised of an algebra of propositions and a comparative confidence relation, $\langle \mathcal{A}, \succsim \rangle$, can be represented in $\langle \mathbb{R}^{\geq 0}, \geq \rangle$ by means of some probability measure. Savage (1954) established sufficient conditions, based on earlier ideas from de Finetti (1931). Kraft, Pratt & Seidenberg (1959) were the first to provide necessary and sufficient conditions for the case of finite algebras, which were presented then in simpler form by Dana Scott (1964).

For the following definition, we take $\boldsymbol{p}^i$ to be the *indicator function* of p. The indicator function of a proposition simply distinguishes those worlds that belong to the proposition from those that don't by assigning 1 to the former and 0 to the latter; i.e., $\boldsymbol{p}^i$ is a function on Ω such that

$$\boldsymbol{p}^i(\omega) = \begin{cases} 1 & \text{if } \omega \in p \\ 0 & \text{otherwise} \end{cases}$$

Definition 4.1 Let $\mathcal{A}$ be an algebra of propositions on Ω, and $\succsim$ a binary relation on $\mathcal{A}$. Then $\langle \mathcal{A}, \succsim \rangle$ is a *finite system of qualitative probability* iff

1. $\mathcal{A}$ is finite (*finitude*)
2. $\succsim$ is complete (*completeness*)
3. $p \succsim \varnothing$ ($\varnothing$*-minimality*)
4. $\Omega \succ \varnothing$ (*non-triviality*)
5. If $p_1, \ldots, p_n$ and $q_1, \ldots, q_n$ are two sequences of propositions in $\mathcal{A}$, then, for $1 \leq j < n$, if
   i) $p_j \succsim q_j$, and
   ii) $\sum_{j=1}^{n} \boldsymbol{p}_j^i(\omega) = \sum_{j=1}^{n} \boldsymbol{q}_j^i(\omega)$ for all $\omega \in \Omega$,

   then $q_n \succsim p_n$ (*Scott's axiom*)

Theorem 4.1 (Scott 1964) *$\langle \mathcal{A}, \succsim \rangle$ is a finite system of qualitative probability iff at least one probability measure μ is a homomorphism from $\langle \mathcal{A}, \succsim \rangle$ into $\langle \mathbb{R}^{\geq 0}, \geq \rangle$.*

Much of the work is done by *Scott's axiom*, but what that axiom says isn't immediately transparent. Very roughly, it tells us that if the two sequences of propositions $p_1, \ldots, p_n$ and $q_1, \ldots, q_n$ contain the same number of truths as a matter of logical necessity, then if the agent is more confident of one proposition in the first sequence than they are of the corresponding proposition in the second sequence, there must be some proposition in the second sequence of which they have more confidence than the corresponding proposition in the first sequence—they must balance out.

But we needn't spend a great deal of time going over what *Scott's axiom* says exactly; more illuminating for present purposes is to consider what the axiom implies in the context of the others. (For extended exposition regarding *Scott's axiom*, see Titelbaum 2022, 491ff.) If we use $\sqcup$ from now on to represent the union of disjoint sets—i.e., the restriction of set-theoretic union $\cup$ to those pairs of sets that have no elements in common—then for any finite system of qualitative probability,

1. $\succsim$ is a weak order (*weak order*)
2. $p \sqcup (q \sqcup r) = (p \sqcup q) \sqcup r$ (*associativity*)
3. $p \sqcup q = q \sqcup p$ (*commutativity*)
4. $p \succsim q$ iff, if $r \cap (p \cup q) = \varnothing$ then $p \sqcup r \succsim q \sqcup r$ ($\sqcup$*-monotonicity*)
5. $p \sqcup q \succsim p$ (*weak positivity*)
6. $p \succsim p \sqcup q$ only if $q \sim \varnothing$ (*minimal identity*)

The *weak order* axiom is, as discussed, necessary for $\succsim$ to be mapped into $\geq$. The *associativity* and *commutativity* axioms fall out of the associativity and commutativity of $\cup$. Finally, *weak positivity* and *minimal identity* correspond to the *non-negativity* condition on the definition of a probability measure, while $\sqcup$*-monotonicity* corresponds to the $\sqcup$*-additivity* condition.

These should hopefully remind you of the properties that permit the additive measurement of length (§2.3), with $\sqcup$ playing something similar to role played by end-to-end concatenation in a positive concatenation structure. Indeed, we can make the similarity clearer by restating Theorem 4.1 thus:

Theorem 4.1′ *$\langle \mathcal{A}, \succsim \rangle$ is a finite system of qualitative probability iff there is at least one probability measure μ that is also a weak homomorphism from $\langle \mathcal{A}, \succsim; \sqcup \rangle$ into $\langle \mathbb{R}^{\geq 0}, \geq; + \rangle$.*

This way of stating Scott's result better captures the *point* of the probabilistic representation of comparative confidence. After all, if the goal was to show how a finite system $\langle \mathcal{A}, \succsim \rangle$ might be represented in $\langle \mathbb{R}^{\geq 0}, \geq \rangle$, then the *weak order* axiom would have sufficed—everything after that just serves to restrict the kinds of qualitative systems under consideration without making any difference to their representability in $\langle \mathbb{R}^{\geq 0}, \geq \rangle$. What makes it worthwhile to represent comparative confidence using a probability measure is that the characteristic properties of such measures—$\sqcup$-*additivity*—correspond to the 'additive' behaviour of $\succsim$ in relation to $\sqcup$. If not for this, then there's no apparent reason to focus on probabilistic representations of $\succsim$ over any number of ordinally-equivalent representations.

With that said, there's a couple important differences with the extensive measurement of length that should be noted. First, additive measures of $\langle \mathbf{L}, \succsim; \circ \rangle$ are *1-point unique*—that is, fixing the numerical value of any non-minimal length L will uniquely determine the remainder of the scale. The same needn't always be true for probabilistic measures of $\langle \mathcal{A}, \succsim; \sqcup \rangle$. Consider a finite algebra with atoms a_1, a_2, a_3, where

$$a_1 \succ (a_2 \cup a_3) \succ a_2 \succ a_3$$

A probability measure μ will represent $\langle \mathcal{A}, \succsim; \sqcup \rangle$ just in case

$$1 > \mu(a_1) > \frac{1}{2} > \mu(a_2 \cup a_3) > \mu(a_2) > \mu(a_3) > 0$$

Obviously, choosing a measure μ such that $\mu(a_1) = \frac{2}{3}$, for instance, won't yet determine the values for a_2 and a_3—it only determines that they'll take distinct positive values summing to $\frac{1}{3}$. Essentially similar examples can be constructed to show that there will be systems of finite probability such that the additive measures of $\langle \mathcal{A}, \succsim; \sqcup \rangle$ are not n-point unique for arbitrarily large n.

Second, $+$ is *meaningful* relative to the additive measures of $\langle \mathbf{L}, \succsim; \circ \rangle$, but the same needn't be true for probabilistic measures of $\langle \mathcal{A}, \succsim; \sqcup \rangle$. In other words, where two μ and μ' are probabilistic representations of the same system $\langle \mathcal{A}, \succsim; \sqcup \rangle$, the relation $R(+, \mu)$ induced on $\mathcal{A}$ by $+$ relative to μ need not be the relation $R(+, \mu')$ induced on $\mathcal{A}$ by $+$ relative to μ'. (Recall from Definition 2.5 that $(p, q, r) \in R(+, \mu)$ iff $\mu(p) + \mu(q) = \mu(r)$.) Consider again the previous example, where $\mu(a_1) = \frac{2}{3}$, and so

$$\mu(a_2 \cup a_3) + \mu(a_2 \cup a_3) = \mu(a_1)$$

Now suppose that μ' is such that $\mu'(a_1) = \frac{3}{4}$; hence

$$\mu'(a_2 \cup a_3) + \mu'(a_2 \cup a_3) \neq \mu'(a_1)$$

Though both μ and μ' are weakly additive measures of $\langle \mathcal{A}, \succsim; \sqcup \rangle$, the qualitative relation corresponding to $+$ under μ isn't identical to the qualitative relation corresponding to $+$ under μ'.

Both disanalogies arise ultimately as a result of the fact that probabilistic representations of a system of qualitative probability need not be unique. This situation can be remedied by the addition of further axioms, such as:

- $p \succsim q$ only if $p \sim q \cup r$ for some $r \in \mathcal{A}$ (*solvability*)

Supposing that $\langle \mathcal{A}, \succsim \rangle$ is a finite system of qualitative probability satisfying *solvability*, then the analogy with the additive measurement of length is considerably stronger. In that case, the set of weakly additive measures of $\langle \mathcal{A}, \succsim; \sqcup \rangle$ will include all and only those φ that are related to μ by a positive similarity transformation, and hence they will be 1-point unique (Suppes 1969, 6–7). Furthermore, $R(+, \mu) = R(+, \varphi)$ for any φ related to μ by a positive similarity transformation, and so $+$ will be meaningful relative to any set of weakly additive measures of $\langle \mathcal{A}, \succsim; \sqcup \rangle$.

Another way to make that analogy clear is to generalise $\sqcup$ slightly, and then show that this generalised relation can be (strongly) mapped into $+$. Start with the following:

Definition 4.2 Where $\sim$ is an equivalence relation and $\bullet$ is a binary operation, $\bullet\backslash\!\sim$ is the *relation induced by* $\bullet$ *and* $\sim$ iff $(p, q, r) \in \bullet\backslash\!\sim$ whenever $p' \bullet q' \sim r$ for some p' and q' such that $p \sim p'$ and $q \sim q'$.

In the special case where $\sim$ is antisymmetric, there's no difference between $\bullet$ and $\bullet\backslash\!\sim$. For example, $+\backslash\!=$ is just the same as $+$. But since the equiprobability of p and q need not imply the identity of p and q, in many cases it will be impossible to construct a system $\langle \mathcal{A}, \succsim; \bullet \rangle$ that admits of an additive measure in the stronger sense. For suppose that $p \bullet q = r$, but there also exists some $s \neq r$ such that $r \sim s$. Then, if φ maps $\succsim$ into $\geq$, then $\bullet$ will map into $+$ only if $\varphi(p) + \varphi(q) = \varphi(s)$ implies $p \bullet q = s$, which by hypothesis is false. But this isn't a deep problem—we dissolve it entirely by mapping the very slightly more general ternary relation $\bullet\backslash\!\sim$ into $+$ instead (where the latter is construed also as a ternary relation). Thus,

Theorem 4.2 *Suppose that* $\langle \mathcal{A}, \succsim \rangle$ *is a finite system of qualitative probability satisfying* solvability*. Then there exists a homomorphism* φ *from* $\langle \mathcal{A}, \succsim, \sqcup\backslash\!\sim \rangle$ *into* $\langle \mathbb{R}^{\geq 0}, \geq, + \rangle$*. Furthermore, the set of all homomorphisms from* $\langle \mathcal{A}, \succsim, \sqcup\backslash\!\sim \rangle$ *into* $\langle \mathbb{R}^{\geq 0}, \geq, + \rangle$ *is unique up to positive similarity transformations, and exactly one of them is a probability measure.*

Proof. Suppose that μ is the unique probability representation of $\langle \mathcal{A}, \succsim \rangle$, guaranteed by the hypothesis of the theorem. $R(+, \mu)$ always maps into $+$ by definition, so for the existence result we need only establish that $R(+, \mu) = \sqcup\backslash\!\sim$. To that end, note that $(p, q, r) \in R(+, \mu)$ iff there exist $p', q', r \in \mathcal{A}$ such that $p' \sim p$, $q' \sim q$, $p' \cap q' = \varnothing$, and $p' \cup q' \sim r$. The right-to-left of that biconditional is trivial, given that $\geq$ represents $\succsim$ and that μ satisfies $\sqcup$-*additivity*. For the left-to-right, suppose $\mu(p) + \mu(q) = \mu(r)$. Where $p \cap q = \varnothing$, let $p = p'$, $q = q'$ and $r = p \cup q$. Where $p \cap q \neq \varnothing$, let s be a proposition such that $s \cap (p \cup q) = \varnothing$ and $s \sim p \cap q$. Using *solvability* it can be shown that some such s exists. Now let $p' = p$, $q' = (q \cup s) - p$, and $r = p \cup q' = p \sqcup q'$. The proof of the uniqueness result is straightforward and omitted. □

Note, thought, that *solvability* isn't necessary for unique probabilistic representation. This is a good thing, since the axiom is very restrictive—in the context of the other axioms, it has the effect of guaranteeing that every atom of $\mathcal{A}$ that's non-minimal in $\succsim$ must be equiprobable with every other such atom; in other words, it forces all non-minimal atoms into a single $\sim$-equivalence class (Suppes 1969, 6–7). A more general condition that also suffices for unique probabilistic representability can be formulated in terms of *scalability*.

Definition 4.3 Suppose $r_1, \ldots, r_n$ is any sequence of pairwise disjoint and equiprobable propositions where $(r_1 \sqcup \ldots \sqcup r_n) \sim q$. Then,

1. If $p \sim r_i$, for $i = 1, \ldots, n$, then p is *directly scaled* by q
2. If p is directly scaled by q, then p is *scaled* by q
3. If p is scaled by q, and q is scaled by r, then p is scaled by r

In other words, the scaling relation is the ancestral of the direct scaling relation. The more general axiom can now be stated with ease:

- For any non-minimal atom $a \in \mathcal{A}$, a is scaled by Ω (*scalability*)

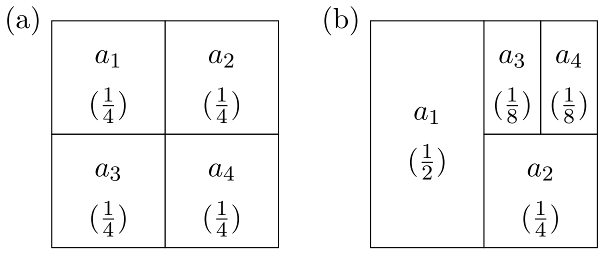

Figure 4.1: *solvability* (a) versus *scalability* (b)

The difference between *solvability* and *scalability* is represented in Figure 4.1. We assume in each case that $\langle \mathcal{A}, \succsim \rangle$ is a finite system of qualitative probability, with the $\succsim$-ordering over the atoms of $\mathcal{A}$ represented by the relative size of the corresponding boxes. On the left, case (a), *solvability* is satisfied, and hence also *scalability*. There are four equiprobable non-minimal atoms, a_1 to a_4, all directly scaled by Ω. Since $\mu(\Omega) = 1$, so each atom must be assigned the value $\frac{1}{4}$ by any probabilistic measure μ. Case (b) violates *solvability*, since

$$(a_1 \cup a_2 \cup a_3) \succsim (a_3 \cup a_4),$$

but there's no p such that

$$(a_1 \cup a_2 \cup a_3) \sim (a_3 \cup a_4 \cup p)$$

However, case (b) still satisfies *scalability*, and is uniquely probabilistically representable. There are four atoms. The largest, a_1, is directly scaled by Ω, since a_1 and $(a_2 \cup a_3 \cup a_4)$ are disjoint, equiprobable, and their union is identical with and thus automatically equiprobable with Ω. So $\mu(a_1) = \frac{1}{2}$ for any probabilistic measure μ. The second largest atom a_2 is then directly scaled by a_1:

$$\frac{1}{2} \cdot \frac{1}{2} = \frac{1}{4}$$

Finally, a_3 and a_4 are both directly scaled by a_2 and thus assigned

$$\frac{1}{4} \cdot \frac{1}{2} = \frac{1}{8}$$

As with *solvability*, *scalability* isn't required for unique probabilistic representability. Necessary and sufficient conditions in terms of $\succsim$ on $\mathcal{A}$ are not easy to express (for reasons explained in Narens 1980), but some will be provided in §5.2 in terms of extended indicator functions.

4.2 The problem of logical omniscience

A probability measure on an algebra of sets $\mathcal{A}$ will always represent a comparative confidence ordering that extends the superset relation over the propositions in $\mathcal{A}$, in the sense that $p \supseteq q$ implies $p \succsim q$. Given **Assumption 1**, Ω includes only logically possible worlds. The combination of these facts presents a problem, since if Ω is restricted to possible worlds then $p \subseteq q$ iff p implies q. In other words, in the presence of **Assumption 1** a probability measure is capable of representing only logically omniscient agents—agents whose comparative confidence orderings invariably respect the logical relations between propositions. Given the desiderata discussed in §3.4, it's therefore worth considering whether and how the standard epistemic approach might be generalised—or better, *de-idealised*—so as to apply also to agents who aren't logically ideal.

The generalisation that I have in mind involves a tweak to the 'concatenation' operation—what we need to do is replace $\sqcup$ with a more general operation that will still allow for the same kind of additivity results that make the standard probabilistic approach interesting, while not forcing logical omniscience. But let me start by noting two important constraints. First, the generalised concatenation operation ought to be *natural*. As noted in §2.4, without naturalness the generalisation task is trivialised and the results of no significance. Second, to avoid disjunctiveness in the theory of belief measurement we are looking for a generalisation of the standard epistemic approach in the sense that we want a qualitative system $\langle \mathcal{A}, \succsim, R \rangle$ that can be represented in $\langle \mathbb{R}^{\geq 0}, \geq, + \rangle$ that includes qualitative probability structures as a special case (which are represented probabilistically), but also allows for the non-probabilistic representation of structures that aren't probabilistically representable. This implies that we need a natural relation R that's an extension of $\sqcup$ in those cases (or at least some of those cases) where $\langle \mathcal{A}, \succsim; \sqcup \rangle$ admits of probabilistic representation.

These are not trivial constraints. It's not so easy to find a natural relation that has this property and which doesn't lead us right back in to the problem of logical omniscience. To appreciate the difficulty here, consider what happens when $R = \sqcup \backslash \sim$. In this case, R is guaranteed to be an extension of $\sqcup$, but mapping $\sqcup \backslash \sim$ into $+$ leads inevitably to logical omniscience. Since p and $\varnothing$ are always disjoint, $p \sqcup \varnothing$ is always defined. Moreover, $\varnothing$ will be the identity element with respect to $\sqcup$ (i.e., for all p, $p \sqcup \varnothing = p$). Consequently, if φ is any additive measure of $\langle \mathcal{A}, \succsim, \sqcup \backslash \sim \rangle$, then $p \sqcup \varnothing = p$ implies $\varphi(p) + \varphi(\varnothing) = \varphi(p)$ implies $\varphi(\varnothing) = 0$. In other words, the identity element of $\sqcup$ will need to be mapped to the identity element of $+$, which is zero. Furthermore, for any $p, q \in \mathcal{A}$, if $q \subseteq p$ then there will exist some $r \in \mathcal{A}$ such that $q \sqcup r = p$, hence $\varphi(q) + \varphi(r) = \varphi(p)$ and so $\varphi(p) \geq \varphi(q)$ and so $p \succsim q$. The result: $p \supseteq q$ implies $p \succsim q$.

To avoid logical omniscience, $\sqcup$ cannot be what "plays the concatenation role". We do better if we consider instead the union of *subjectively incompatible* propositions. Henceforth, let $\uplus$ designate this operation, defined relative to $\succsim$ as the restriction of $\cup$ to those pairs of propositions p, q such that $p \cap q$ is minimal in $\succsim$. Intuitively, p and q are subjectively incompatible whenever the subject has at least as much confidence any proposition whatsoever as they do in the conjunction of p and q. Then, $\uplus$ is an extension of $\sqcup$ at least whenever $\langle \mathcal{A}, \succsim; \sqcup \rangle$ can be represented probabilistically. In other cases, $p \sqcup q = r$ needn't imply $p \uplus q = r$. Hence, it's possible to have an additive mapping φ from $\langle \mathcal{A}, \succsim, \uplus \backslash \sim \rangle$ into $\langle \mathbb{R}^{\geq 0}, \geq, + \rangle$ such that φ needn't satisfy $\sqcup$-*additivity* in general, but is also guaranteed to satisfy $\sqcup$-*additivity* in all cases where a probability representation of $\langle \mathcal{A}, \succsim, \uplus \backslash \sim \rangle$ exists.[1]

I'll start with a simple example, chosen to demonstrate that none of $\varnothing$-*minimality*, *non-triviality* or *Scott's axiom* are required for the specified homomorphisms to exist. (As such, it's intended to be an extreme example, not a realistic one.) We suppose that $\mathcal{A}$ contains exactly four atoms, a_1 through a_4. We then label the non-atomic propositions using the indices of the atomic propositions from which they're constructed; so, for instance,

$$p_{\langle 234 \rangle} = a_2 \cup a_3 \cup a_4$$

Given that, consider the following extremely non-omniscient confidence ranking:

$$\begin{bmatrix} \Omega \\ \varnothing \end{bmatrix} \succ \begin{bmatrix} a_1 \\ a_2 \\ a_4 \\ p_{\langle 134 \rangle} \\ p_{\langle 234 \rangle} \end{bmatrix} \succ \begin{bmatrix} p_{\langle 12 \rangle} \\ p_{\langle 14 \rangle} \\ p_{\langle 24 \rangle} \\ p_{\langle 34 \rangle} \\ p_{\langle 123 \rangle} \end{bmatrix} \succ \begin{bmatrix} p_{\langle 13 \rangle} \\ p_{\langle 23 \rangle} \\ p_{\langle 124 \rangle} \end{bmatrix} \succ a_3$$

We want to show that there is at least one function $\varphi : \mathcal{A} \mapsto \mathbb{R}^{\geq 0}$ such that for all propositions $p, q \in \mathcal{A}$,

i $p \succsim q$ iff $\varphi(p) \geq \varphi(q)$
ii $(p, q, r) \in \uplus \backslash \sim$ iff $\varphi(p) + \varphi(q) = \varphi(r)$

It's clear that some φ satisfies the first property—e.g.,

$$\varphi(\Omega) = 1, \quad \varphi(a_1) = \frac{3}{4}, \quad \varphi(p_{\langle 12 \rangle}) = \frac{1}{2}, \quad \varphi(p_{\langle 13 \rangle}) = \frac{1}{4}, \quad \varphi(a_3) = 0$$

[1] Every homomorphic mapping from $\langle \mathcal{A}, \succsim, \uplus \backslash \sim \rangle$ into $\langle \mathbb{R}^{\geq 0}, \geq, + \rangle$ is a *weakly* additive measure of $\langle \mathcal{A}, \succsim; \uplus \rangle$. For the reasons discussed earlier, strongly additive measures of $\langle \mathcal{A}, \succsim; \uplus \rangle$ will often be impossible where $\succsim$ isn't antisymmetric.

So we just need to show that φ defined as such satisfies property ii. To that end, note that p and q are subjectively incompatible only if they both include the minimal proposition a_3; for all the other propositions, it matters not where they sit in the $\succsim$ ordering so long as they don't sit at the bottom. Hence, we need only consider the *concatenable* propositions:

$$\Omega \succ \begin{bmatrix} p_{\langle 134 \rangle} \\ p_{\langle 234 \rangle} \end{bmatrix} \succ \begin{bmatrix} p_{\langle 34 \rangle} \\ p_{\langle 123 \rangle} \end{bmatrix} \succ \begin{bmatrix} p_{\langle 13 \rangle} \\ p_{\langle 23 \rangle} \end{bmatrix} \succ a_3$$

It's then easy to check that $\varphi(p \uplus q) = \varphi(p) + \varphi(q)$; and whenever $\varphi(p) + \varphi(q) = \varphi(r)$, then $(p, q, r) \in \uplus \backslash \sim$. Thus, it's possible to have a weakly additive (but not $\sqcup$-*additive*) measure of $\langle \mathcal{A}, \succsim; \uplus \rangle$. More generally, it's possible to have a (strong) homomorphism from $\langle \mathcal{A}, \succsim, \uplus \backslash \sim \rangle$ into $\langle \mathbb{R}^{\geq 0}, \geq, + \rangle$, even while $\succsim$ is quite far from logically omniscient.

The construction makes use of the same general notion of *scaling* that was discussed in the previous section, though this time understood in terms of pairwise subjectively incompatible propositions rather than pairwise disjoint propositions. For example, $p_{\langle 34 \rangle}$ and $p_{\langle 123 \rangle}$ are equiprobable and subjectively incompatible, and their union is Ω; hence, they're directly scaled by Ω:

$$\varphi(p_{\langle 34 \rangle}) = \varphi(p_{\langle 123 \rangle}) = \frac{1}{2}\varphi(\Omega)$$

Then, $p_{\langle 13 \rangle}$ and $p_{\langle 23 \rangle}$ are equiprobable and subjectively incompatible, and their union is $p_{\langle 123 \rangle}$; hence, they're scaled by $p_{\langle 123 \rangle}$ and derivatively scaled by Ω:

$$\varphi(p_{\langle 13 \rangle}) = \varphi(p_{\langle 123 \rangle}) = \frac{1}{4}\varphi(\Omega)$$

The value for $p_{\langle 134 \rangle}$ can then be determined by summing the values for the subjectively incompatible propositions $p_{\langle 13 \rangle}$ and $p_{\langle 34 \rangle}$; i.e.,

$$\frac{1}{4}\varphi(\Omega) + \frac{1}{2}\varphi(\Omega) = \frac{3}{4}\varphi(\Omega)$$

Similar applies to $p_{\langle 234 \rangle}$. And the, finally, the value for every other proposition is determined via equiprobability with some proposition whose value has already been determined in via scaling relative to Ω. What makes it possible for this construction to avoid logical omniscience is that subjective incompatability needn't coincide with logical incompatible. It *will* so coincide whenever $\langle \mathcal{A}, \succsim; \sqcup \rangle$ can be represented probabilistically; thus we can generalise the case of probabilistic representations by swapping out $\sqcup$ as the concatenation operation for the more general $\uplus$ operation.

Sufficient conditions for the existence of such representations are established in the following definition and associated theorem. There are three structural conditions—*finitude*, *richness* and *weak solvability*—all of which are satisfied in the foregoing example.

Definition 4.4 Let $\mathcal{A}$ be an algebra of propositions on Ω, and $\succsim$ a binary relation on $\mathcal{A}$. Then $\langle \mathcal{A}, \succsim \rangle$ is a *finite system of additive confidence* iff $\mathcal{A}$ is finite and for all $p, q, r, s \in \mathcal{A}$,

1. $\succsim$ is a weak order (*weak order*)
2. If p and q are subjectively incompatible, $p \succsim r$ and $q \succsim s$, then there are r' and s' such that $r \sim r'$, $s \sim s'$, and r' and s' are subjectively incompatible (*richness*)
3. If $p \succ q$, then there are subjectively incompatible q' and r such that $q' \sim q$ and $p \succsim q' \cup r$ (*weak solvability*)
4. If p and r are subjectively incompatible, q and r are subjectively incompatible, and $p \succsim q$, then $p \cup r \succsim q \cup r$ ($\uplus$-*monotonicity*)
5. If p and q are subjectively incompatible, then $p \cup q \succsim p$, with $p \succsim p \cup q$ only if q is minimal ($\uplus$-*positivity*)

Theorem 4.3 *If $\langle \mathcal{A}, \succsim \rangle$ is a finite system of additive confidence, then there exists a homomorphism from $\langle \mathcal{A}, \succsim, \uplus \backslash \sim \rangle$ into $\langle \mathbb{R}^{\geq 0}, \geq, + \rangle$; furthermore, the set of all such homomorphisms is unique up to positive similarity transformations.*

Proof. The finer details of the proof are not especially illuminating, so I provide instead a summary. The strategy is to reconstruct $\langle \mathcal{A}, \succsim, \uplus \backslash \sim \rangle$ as a system for which strongly additive measures are known to exist. First we let $\mathbf{A} = \{\mathbf{p}, \mathbf{q}, \dots\}$ be the set of $\sim$-equivalence classes in $\mathcal{A}$, with the minimal elements excised; that is, $\mathbf{p} = \{q \in \mathcal{A} \,|\, q \sim p\}$, with $\mathbf{p} \in \mathbf{A}$ only if $p \succ \varnothing$. We then let $\boldsymbol{\succsim}$ be the total order induced on $\mathbf{A}$ by $\succsim$; that is, $\mathbf{p} \boldsymbol{\succsim} \mathbf{q}$ whenever $p \succsim q$. C is to be interpreted as the set of concatenable pairs in $\mathbf{A}$, so $(\mathbf{p}, \mathbf{q}) \in C$ just when $p' \uplus q'$ is defined for some $p' \in \mathbf{p}$ and $q' \in \mathbf{q}$, or (same thing) when $(p, q, r) \in \uplus \backslash \sim$ for some r. Finally, $\circ$ is an operation on $\mathbf{A}$ such that $\mathbf{p} \circ \mathbf{q} = \mathbf{r}$ iff $(p, q, r) \in \uplus \backslash \sim$, and so a function from C into $\mathbf{A}$. We then want to show that $\langle \mathbf{A}, \boldsymbol{\succsim}, C; \circ \rangle$ satisfies:

A. $\boldsymbol{\succsim}$ is a total order
B. If $(\mathbf{p}, \mathbf{q}) \in C$, $\mathbf{p} \boldsymbol{\succsim} \mathbf{r}$, and $\mathbf{q} \boldsymbol{\succsim} \mathbf{s}$, then $(\mathbf{r}, \mathbf{s}) \in C$
C. If $(\mathbf{r}, \mathbf{p}) \in C$, then if $\mathbf{p} \boldsymbol{\succsim} \mathbf{q}$, $\mathbf{r} \circ \mathbf{p} \boldsymbol{\succsim} \mathbf{r} \circ \mathbf{q}$
D. If $(\mathbf{p}, \mathbf{r}) \in C$, then if $\mathbf{p} \boldsymbol{\succsim} \mathbf{q}$, $\mathbf{p} \circ \mathbf{r} \boldsymbol{\succsim} \mathbf{q} \circ \mathbf{r}$
E. $(\mathbf{p}, \mathbf{q}), (\mathbf{p} \circ \mathbf{q}, \mathbf{r}) \in C$ iff $(\mathbf{q}, \mathbf{r}), (\mathbf{p}, \mathbf{q} \circ \mathbf{r}) \in C$, and when both hold then $(\mathbf{p} \circ \mathbf{q}) \circ \mathbf{r} = \mathbf{p} \circ (\mathbf{q} \circ \mathbf{r})$
F. If $(\mathbf{p}, \mathbf{q}) \in C$, then $\mathbf{p} \circ \mathbf{q} \boldsymbol{\succ} \mathbf{p}$
G. If $\mathbf{p} \boldsymbol{\succ} \mathbf{q}$, then there exists an $\mathbf{r} \in \mathbf{A}$ such that $(\mathbf{q}, \mathbf{r}) \in C$ and $\mathbf{p} \boldsymbol{\succsim} \mathbf{q} \circ \mathbf{r}$

A follows from *weak order*, and B from *richness*. Given B, conditions C and D follow from $\uplus$-*monotonicity* and the commutativity of $\cap$ and $\cup$. The first conjunct of E falls out of how $\circ$ has been defined, and the second conjunct follows from the associativity of $\cap$ and $\cup$. F is fixed by $\uplus$-*positivity*, and G by *weak solvability*. From these seven conditions plus *finitude*, it follows that the system $\langle \mathbf{A}, \boldsymbol{\succsim}, C; \circ \rangle$ is a *Archimedean, regular, positive, ordered local semigroup* (Krantz *et al.* 1971, 44–5). This suffices for the existence of a homomorphism ψ from $\langle \mathbf{A}, \boldsymbol{\succsim}; \circ \rangle$ into $\langle \mathbb{R}^{>0}, \geq; + \rangle$, and the set of such homomorphisms is unique up to positive similarity transformations. (This is a corollary of Krantz *et al.* 1971, 44–6, Theorem 4 and Theorem 4′.) We then let φ be defined on $\mathcal{A}$ such that $\varphi(p) = \psi(\mathbf{p})$ for all non-minimal p, and $\varphi(p) = 0$ otherwise, which gives us a homomorphism from $\langle \mathcal{A}, \succsim, \uplus \backslash \sim \rangle$ into $\langle \mathbb{R}^{\geq 0}, \geq, + \rangle$, and inherits the uniqueness properties mentioned above. □

In the special case where $\langle \mathcal{A}, \succsim; \sqcup \rangle$ is a finite system of qualitative probability satisfying *solvability*, then it will also be a finite system of additive confidence—and in that case, the unique representation φ of $\langle \mathcal{A}, \succsim, \uplus \backslash \sim \rangle$ in $\langle \mathbb{R}^{\geq 0}, \geq, + \rangle$ that satisfies *normalisation* will also be the unique probability measure that represents $\langle \mathcal{A}, \succsim; \sqcup \rangle$. Thus, Theorem 4.4 is a generalisation of Theorem 4.2, with strictly weaker axioms on $\succsim$ and, as a result, a strictly more flexible form of non-necessarily-probabilistic representation that doesn't imply logical omniscience.

However, it's not all good news. Theorem 4.4 offers a step forward in dealing with the problem of logical omniscience, though no great leap. While we've managed to avoid the strict form of logical omniscience—i.e., where $p \supseteq q$ always implies $p \succsim q$—the additive representation of $\langle \mathcal{A}, \succsim, \uplus \backslash \sim \rangle$ is perhaps not as flexible as one might like for a theory of belief measurement applicable to logically non-ideal agents. For one thing, note that Ω will always be maximal in $\succsim$. To see why, suppose it isn't. A proposition p is *concatenable* just in case it's a superset of q for some q that's minimal in $\succsim$. The concatenable propositions are those that can stand in relations of subjective incompatibility. In a finite system of additive confidence, every proposition will be equiprobable with a concatenable proposition. So, if Ω isn't maximal in $\succsim$, then at least one other concatenable proposition must be. Let p_{max} be that proposition, or one of them, and let p_{min} be any minimal proposition that implies p_{max}. Now suppose that q is $(\Omega \setminus p_{max}) \cup p_{min}$. So q and p_{max} are subjectively incompatible, and we should have $\varphi(p) + \varphi(q) = \varphi(p \cup q)$; but $p \cup q = \Omega$, so $\varphi(\Omega) \geq \varphi(p)$, contradicting the hypothesis that $p_{max} \succ \Omega$.

More generally, it will be noted that in a finite system of additive confidence, $q \subseteq p$ will still imply $p \succsim q$ *with respect to* concatenable propositions p and q. It's for this reason, ultimately, that Ω must always sit at the top of the $\succsim$ ranking, since Ω is an *absorbing element* with respect to $\uplus$ (i.e., $\Omega \uplus p = \Omega$ for any p). This also means that whenever $\varnothing$ is minimal, then the stricter form of logical omniscience follows immediately—since in that case every proposition in the algebra is automatically concatenable.

The upshot is that, while it's possible to maintain the analogy with the measurement of length while avoiding logical omniscience, the results here are still quite limited. It's not logical omniscience *per se*, but it's not that far from it either. Other generalisations *might* be possible, though the relevant work has yet to be done. The hard part, as I said, is finding an appropriately natural operation to "play the concatenation role", which covers both the probabilistic and non-probabilistic case and doesn't lead us right back to logical omniscience or something near as bad.

Chapter 5

Epistemic Approaches: Alternatives

Epistemic approaches to the measurement of belief aren't limited to those involving a single complete binary comparative confidence relation. In this chapter, I briefly survey a number of alternative epistemic approaches. The first involves quarternary conditional confidence relations (§5.2); then qualitative expectation relations (§5.1); then structures composed out of multiple doxastic relations (§5.3).

5.1 Conditional confidence

A theory of belief measurement that employs a binary confidence relation will be suited for numerical representations that assign a single numerical value to each proposition intended to represent the agent's *unconditional* degree of confidence towards that proposition. However, it's sometimes thought that the more fundamental concept in epistemology is not unconditional confidence but rather conditional confidence—the level of confidence one has in p *given* a hypothesis q (e.g., Hájek 2003). A common motivation for this is that, while it's standard to define conditional probabilities out of unconditional probabilities like so

$$\mu(p|q) = \frac{\mu(p \cap q)}{\mu(q)},$$

that definition only makes sense when $\mu(q) > 0$; yet there appear to be cases where it makes sense to speak of the probability of p conditional on q even while the unconditional probability of q is zero.

An epistemic approach to the measurement of belief that fits with this perspective replaces the binary confidence relation on $\mathcal{A}$ with a quarternary confidence relation instead—or, same thing, a binary relation $\succsim$ on $\mathcal{A}\times\mathcal{A}$, interpreted

$$(p, q) \succsim (r, s) \text{ iff } \alpha \text{ is at least as confident in } p \text{ given } q \text{ as she is in } r \text{ given } s$$

To make things a little easier, let's write $p|q \succsim r|s$ instead. The idea, then is to lay down axioms on $\succsim$ that will suffice for the 'probabilistic' representation thereof.

Much of the work done on this matter is owing to Koopman (see especially his 1940a and 1940b), though for this section I briefly summarise a more recent result in (Hawthorne 2016).[1] Since we are taking conditional probabilities as basic, the numerical representation cannot consist in probability measures as per Definition 3.2. Instead, we employ *Popper functions*, which generalise the classic definition of a probability measure:

Definition 5.1 $\pi : \mathcal{A} \times \mathcal{A} \mapsto \mathbb{R}$ is a *Popper function* iff

1. For some $p, q, r, s \in \mathcal{A}$, $\pi(p|q) \neq \pi(r|s)$
2. For all $p, q, r \in \mathcal{A}$, $\pi(p|p) \geq \pi(q|r)$
3. If $q \subseteq p$, then $\pi(p|r) \geq \pi(q|r)$
4. $\pi(p|q) + \pi(\neg p|q) = \pi(q|q)$ unless $\pi(r|q) = \pi(q|q)$ for all $r \in \mathcal{A}$
5. $\pi(p \cap q|r) = \pi(p|q \cap r) \times \pi(q|r)$

Relative to a fixed condition, a Popper function behaves essentially like a probability measure. For instance, fixing the condition to Ω, the definition implies:

- $\pi(p|\Omega) \in [0, 1]$,
- $\pi(\Omega|\Omega) = 1$, and
- if $p \cap q = \varnothing$, then $\pi(p \cup q|\Omega) = \pi(p|\Omega) + \pi(q|\Omega)$

Moreover, if μ is the probability measure corresponding to $\pi(\cdot|\Omega)$, then for any p such that $\pi(p|\Omega) > 0$, $\pi(q|p)$ will behave just like $\mu(q|p)$. The difference, though, is that $\pi(q|p)$ can still be defined even when $\pi(p|\Omega) = 0$, whereas $\mu(q|p)$ cannot be defined in the ordinary fashion when $\mu(p) = 0$. In this case, $\pi(\cdot|p)$ *also* behaves just like a probability measure μ', different from μ, in the same way that $\pi(\cdot|\Omega)$ behaves like μ. And likewise, there may be some r such that $\pi(r|p) = 0$, and $\pi(\cdot|r)$ might behave in turn like yet another probability measure different again from μ'. Thus the Popper function π can act like an ordered hierarchy of probability measures. As Hawthorne helpfully puts it,

> … a Popper function may consist of a ranked hierarchy of classical probability functions, where conditionalization on a probability 0 sentence induces a transition from one classical probability function to another classical function at a lower rank. The idea is that probability 0 need not mean "absolutely impossible". Rather, it means something like, "not a viable possibility unless (and until) the more plausible alternatives are refuted." (2016, 281)

See also (Halpern 2001) and (Brickhill & Horsten 2018) for detailed discussion on the close relationship between Popper functions, lexicographic probability measures (lexically-ordered sequences of probabilities), and non-Archimedean probability measures (probabilities that can take infinitesimal values).

[1] Hawthorne suggests an interpretation of $\succsim$ as a relation of comparative evidential support between premises and conclusions, not—or not necessarily—as a comparative confidence relation. As he notes, though, the formalism can be interpreted in a variety of ways. See (DiBella 2018) for an instance of the quarternary $\succsim$ explicitly interpreted as a relation of comparative conditional confidence.

As one might expect, the extra complexity of the numerical representation—with both an additive component in condition 4 multiplicative component in condition 5—corresponds to significant increased complexity in the required axioms on $\succsim$:

Definition 5.2 Let $\mathcal{A}$ be an algebra of propositions on Ω, and $\succsim$ a binary relation on $\mathcal{A}\times\mathcal{A}$. We say that $\langle \mathcal{A}\times\mathcal{A}, \succsim\rangle$ is a *system of qualitative conditional probability* iff the following are satisfied:

1. $\succsim$ is a weak order (*weak order*)
2. For some $p, q, r, s \in \mathcal{A}$, $p|q \succ r|s$ (*non-triviality*)
3. For all $p, q \in \mathcal{A}$, $p|p \succsim q|p$ (*maximality*)
4. For all $p, q, r \in \mathcal{A}$, if $p \subseteq q$, then $q|r \succsim p|r$ (*implication*)
5. For all $p, q, r, s \in \mathcal{A}$, if $p|q \succsim r|s$ and $q \neq \varnothing$, then $\neg r|s \succsim \neg p|q$ (*negation-symmetry*)
6. For all $p_1, q_1, r_1, p_2, q_2, r_2 \in \mathcal{A}$, if
   i) $p_1|(q_1 \cap r_1) \succsim p_2|(q_2 \cap r_2)$ and $q_1|r_1 \succsim q_2|r_2$, or
   ii) $p_1|(q_1 \cap r_1) \succsim q_2|r_2$ and $q_1|r_1 \succsim p_2|(q_2 \cap r_2)$,

   then $(p_1 \cap q_1)|r_1 \succsim (p_2 \cap q_2)|r_2$ (*composition*)
7. For all $p_1, q_1, r_1, p_2, q_2, r_2 \in \mathcal{A}$, if $(p_1 \cap q_1)|r_1 \succsim (p_2 \cap q_2)|r_2$ and $r_2 \not\subseteq \neg q_2$, and
   i) if $q_2|r_2 \succsim q_1|r_1$, then $p_1|(q_1 \cap r_1) \succsim p_2|(q_2 \cap r_2)$
   ii) if $q_2|r_2 \succsim p_1|(q_1 \cap r_1)$, then $q_1|r_1 \succsim p_2|(q_2 \cap r_2)$ (*decomposition-a*)
8. For all $p_1, q_1, r_1, p_2, q_2, r_2 \in \mathcal{A}$, if $(p_1 \cap q_1)|r_1 \succsim (p_2 \cap q_2)|r_2$ and $(q_2 \cap r_2) \not\subseteq \neg p_2$, then
   i) if $p_2|(q_2 \cap r_2) \succsim p_1|(q_1 \cap r_1)$, then $q_1|r_1 \succsim q_2|r_2$
   ii) if $p_2|(q_2 \cap r_2) \succsim q_1|r_1$, then $p_1|(q_1 \cap r_1) \succsim q_2|r_2$ (*decomposition-b*)
9. For all $p, q, r, s \in \mathcal{A}$, if $p|q \succ r|s$, then for some $n \geq 2$ there exist $t_1, \ldots, t_n, u \in \mathcal{A}$ such that
   i) $u|u \succ \neg t_1|u$,
   ii) for distinct $i, j = 1, \ldots, n$, $t_i|u \sim t_j|u$ and $\neg(t_i \cap t_j)|u \succsim u|u$,
   iii) $(t_1 \cup \ldots \cup t_n)|u \succsim u|u$,
   iv) for some $m \leq n$, $p|q \succ (t_1 \cup \ldots \cup t_m)|u \succ r|s$ (*Archimedean*)

Theorem 5.1 (Hawthorne 2016) *If $\langle \mathcal{A}\times\mathcal{A}, \succsim\rangle$ is a system of qualitative conditional probability, there exists a homomorphism from $\langle \mathcal{A}\times\mathcal{A}, \succsim\rangle$ into $\langle \mathbb{R}^{\geq 0}, \geq\rangle$, and exactly one such homomorphism is a Popper function.*

The *non-triviality*, *maximality* and *implication* axioms directly correspond to conditions 1, 2 and 3 of Definition 5.1. The *negation-symmetry* axiom is the main axiom corresponding to the additivity condition 4, while the *composition* and *decomposition* axioms correspond to the multiplicative condition. The *Archimedean* axiom says that whenever $p|q \succ r|s$, there is a finite number of mutually exclusive and equiprobable propositions such that the conditional probability of their union (relative to some condition) is strictly between that of $p|q$ and $r|s$. In terms of the representation: if $p|q \succ r|s$ then the difference between $\pi(p|q)$ and $\pi(r|s)$ is not infinitesimal, ensuring that $\succsim$ can be represented in $\mathbb{R}$.

5.2 Qualitative expectations

A rather different epistemic approach—originating with Suppes & Zanotti (1976), see also Clark (2000) and (Suppes & Pedersen 2016)—takes the primitive ordering relation $\succsim$ to be defined not over an algebra of propositions, but instead over an algebra of *extended indicator functions*.

Extended indicator functions are a generalisation of indicator functions. In the broadest terms, an extended indicator function is a certain kind of random variable—an integer-valued function f defined on Ω such that for some positive integer n, propositions $p_1, \ldots, p_n$, and non-negative integers $k_1, \ldots, k_n$,

$$f(\omega) = \sum_{j=1}^{n} k_j \cdot \boldsymbol{p}_j^i(\omega)$$

But that's unlikely to be intuitive, so it'll help to consider how extended indicator functions can be built up via the pointwise summation of ordinary indicator functions. Start with the indicator function of p, or $\boldsymbol{p}^i$, which in Chapter 4 was defined as the function that takes each world ω in Ω and returns the value 1 if ω belongs to p, and 0 otherwise. Now consider it's n^{th} iteration, $n\boldsymbol{p}^i$, defined:

$$n\boldsymbol{p}^i(\omega) \;=\; \overbrace{\boldsymbol{p}^i(\omega) + \cdots + \boldsymbol{p}^i(\omega)}^{n \text{ times}} \;=\; \begin{cases} n & \text{if } \omega \in p \\ 0 & \text{otherwise} \end{cases}$$

For any integer $n \geq 1$, the n^{th} iteration of any indicator function will count as an extended indicator function. Clearly, where $n = 1$, then $1\boldsymbol{p}^i = \boldsymbol{p}^i$; and where $n > 1$, then $n\boldsymbol{p}^i$ can be expressed as the pointwise sum of $m\boldsymbol{p}^i$ and $k\boldsymbol{p}^i$ (symbolised $m\boldsymbol{p}^i \dotplus k\boldsymbol{p}^i$) for $m + k = n$. More generally, the pointwise sum of any two extended indicator functions will count as an extended indicator function. So, for example, $n\boldsymbol{p}^i \dotplus m\boldsymbol{q}^i$ is an extended indicator function:

$$n\boldsymbol{p}^i \dotplus m\boldsymbol{q}^i(\omega) \;=\; n\boldsymbol{p}^i(\omega) + m\boldsymbol{q}^i(\omega) \;=\; \begin{cases} n + m & \text{if } \omega \in p \text{ and } \omega \in q \\ n & \text{if } \omega \in p \text{ and } \omega \notin q \\ m & \text{if } \omega \notin p \text{ and } \omega \in q \\ 0 & \text{otherwise} \end{cases}$$

In the same fashion, $(n\boldsymbol{p}^i \dotplus m\boldsymbol{q}^i) \dotplus k\boldsymbol{r}^i$ is an extended indicator function, and so on. Hence we can construct a space of extended indicator functions by starting with a set of propositions, taking the set of indicator functions corresponding to those propositions, and closing it under pointwise summation:

Definition 5.3 $\boldsymbol{\mathcal{A}}^i$ is the *algebra of extended indicator functions generated by* $\mathcal{A}$ iff

1. For all $p \in \mathcal{A}$, $\boldsymbol{p}^i \in \boldsymbol{\mathcal{A}}^i$
2. If $f, g \in \boldsymbol{\mathcal{A}}^i$, then $f \dotplus g \in \boldsymbol{\mathcal{A}}^i$
3. Nothing else is in $\boldsymbol{\mathcal{A}}^i$

This algebra of extended indicator functions will comprise the domain of the primitive binary relation $\succsim$—a *qualitative expectations* relation—with the goal being to represent $\succsim$ via an *expectation function*:

Definition 5.4 Where $\boldsymbol{\mathcal{A}}^i$ is the algebra of extended indicator functions generated by $\mathcal{A}$, a function $\epsilon : \boldsymbol{\mathcal{A}}^i \mapsto \mathbb{R}^{\geq 0}$ is an *expectation function* iff for all $x, y \in \boldsymbol{\mathcal{A}}^i$,

1. $\epsilon(\boldsymbol{\Omega}^i) > \epsilon(\boldsymbol{\varnothing}^i) = 0$,
2. $\epsilon(x \dotplus y) = \epsilon(x) + \epsilon(y)$

Sufficient conditions for the existence of such representations are provided in the following theorem. Given the additive structure of the desired representation, these should come as no surprise:

Theorem 5.2 (Suppes 2016) *Let $\mathcal{A}$ be an algebra of propositions on Ω, and $\succsim$ a binary relation on the algebra $\boldsymbol{\mathcal{A}}^i$ of extended indicator functions generated by $\mathcal{A}$. Suppose that $\langle \boldsymbol{\mathcal{A}}^i, \succsim \rangle$ satisfies the following, for all for all $x, y, z \in \boldsymbol{\mathcal{A}}^i$,*

1. *$\succsim$ is a weak order* *(weak order)*
2. *$x \succsim \boldsymbol{\varnothing}^i$* *($\boldsymbol{\varnothing}^i$-minimality)*
3. *$\boldsymbol{\Omega}^i \succ \boldsymbol{\varnothing}^i$* *(non-triviality)*
4. *$x \succsim y$ iff $x \dotplus z \succsim y \dotplus z$* *($\dotplus$-monotonicity)*
5. *If $x \succ y$, then there are $k, n \geq 1$ with $nx \succ k\boldsymbol{\Omega}^i \succ ny$* *(Archimedean)*

Then at least one expectation function is a homomorphism from $\langle \boldsymbol{\mathcal{A}}^i, \succsim \rangle$ into $\langle \mathbb{R}^{\geq 0}, \geq \rangle$; furthermore, the set of homomorphisms from $\langle \boldsymbol{\mathcal{A}}^i, \succsim \rangle$ into $\langle \mathbb{R}^{\geq 0}, \geq \rangle$ that are also expectation functions is unique up to positive similarity transformations.

Note that any expectation function that maps $\langle \boldsymbol{\mathcal{A}}^i, \succsim \rangle$ into $\langle \mathbb{R}^{\geq 0}, \geq \rangle$ is *ipso facto* a weakly additive representation of $\langle \boldsymbol{\mathcal{A}}^i, \succsim; \dotplus \rangle$ in $\langle \mathbb{R}^{\geq 0}, \geq; + \rangle$, and vice versa. Indeed, similar to the reformulation Theorem 4.1 earlier, it's possible to re-write Theorem 5.2 so as to make the underlying connection with extensive measurement a bit more apparent. Essentially: if $\langle \boldsymbol{\mathcal{A}}^i, \succsim; \dotplus \rangle$ satisfies the stated axioms, then there is a weak homomorphism from $\langle \boldsymbol{\mathcal{A}}^i, \succsim; \dotplus \rangle$ into $\langle \mathbb{R}^{\geq 0}, \geq; + \rangle$; and furthermore, the set of all such homomorphisms is unique up to positive similarity transformations.

Furthermore, note that any expectation function ϵ is related by a positive similarity transformation to exactly one normalised expectation function ϵ', with $\epsilon'(\boldsymbol{\Omega}^i) = 1$. This ϵ' defines a probability measure μ if, for all $p \in \mathcal{A}$, we let $\mu(p) = \epsilon'(\boldsymbol{p}^i)$. In other words, corresponding to the set of weakly additive measures of $\langle \boldsymbol{\mathcal{A}}^i, \succsim; \dotplus \rangle$ is a unique probability measure on $\mathcal{A}$. So, clearly, there is a close connection between expectation representations of qualitative expectation relations and probabilistic representations of comparative confidence relations. Indeed, say that $\succsim$ on $\mathcal{A}$ is the weak order induced by $\succsim$ on $\boldsymbol{\mathcal{A}}^i$ whenever

$$p \succsim q \text{ iff } \boldsymbol{p}^i \succsim \boldsymbol{q}^i$$

Then, Suppes & Zanotti (1976, 435–7) were able to establish that $\langle \mathcal{A}, \succsim \rangle$ has a *unique* probabilistic representation if and only if there exists a system of qualitative expectations $\langle \boldsymbol{\mathcal{A}}^i, \succsim \rangle$ such that $\succsim$ on $\mathcal{A}$ is the weak order induced by $\succsim$ on $\boldsymbol{\mathcal{A}}^i$. To put that another way: a (complete) binary confidence relation is uniquely probabilistically representable just when it can be extended to a qualitative expectations relation which satisfies Suppes & Zanotti's five axioms.

So much for the formalities, now for the hard part: the interpretation of the *qualitative expectation* relation isn't straightforward, and I suspect that this is main reason why there's been comparatively little work done on this approach. In the usual case, random variables are functions from the outcomes of a experiment-type to numerical values of those outcomes. For instance, if we say the experiment is tossing two six-sided die, there are 36 possible outcomes corresponding to the different combinations, and 11 possible numerical values from 2 to 12 they might sum to. Letting $\mathbf{r}$ be the corresponding random variable, the expected value ϵ of $\mathbf{r}$ is the probability-weighted average value of the outcomes (under the supposition the experiment is run), and the sum of the expected value of $\mathbf{r}$ with itself n times can be interpreted as the expected total value of n independent runs of the same experiment under the same conditions. If the die are fair, then $\epsilon(\mathbf{r}) = 7$, and

$$\epsilon(\mathbf{r} \dotplus \mathbf{r}) = \epsilon(\mathbf{r}) + \epsilon(\mathbf{r}) = 14$$

For this to make sense, though, it must be possible for those 36 outcomes to recur across independent instances of the same experiment. It is much less clear how to make sense of the iterated variables where the 'outcomes' are maximally specific possible worlds and the 'experiment', as such, can only be run once. Suppose p is the proposition *there are dogs*, and q the proposition *most roses are red*. Presumably we should be able to find both in $\mathcal{A}$ given the intended interpretation. Each corresponds very naturally to a random variable over Ω, namely $\boldsymbol{p}^i$ and $\boldsymbol{q}^i$, and there's no difficulty in interpreting $\boldsymbol{p}^i \succsim \boldsymbol{q}^i$ as an expectation relation *in this case*. The interpretation of $3\boldsymbol{p}^i$ and $5\boldsymbol{q}^i$ is not at all transparent, by contrast, and still less the interpretation of $3\boldsymbol{p}^i \dotplus 5\boldsymbol{q}^i$.

In connection to this, it's noteworthy that Suppes (2014, 53) later flagged interpretive difficulties as a distinctive cost for the approach, particularly vis-à-vis the reading of mixed indicator functions $\boldsymbol{p}^i_j \dotplus \boldsymbol{q}^i_j$. Suppes & Zanotti explain one possible way to interpret their mixed non-iterated functions $\boldsymbol{p}^i \dotplus \boldsymbol{q}^i$ thus:

> Suppose Smith is considering two locations to fly to for a weekend vacation. Let p_j be the event of sunny weather at location j and q_j be the event of warm weather at location j. The qualitative comparison Smith is interested in is the expected value of $\boldsymbol{p}^i_1 \dotplus \boldsymbol{q}^i_1$ *versus* the expected value of $\boldsymbol{p}^i_2 \dotplus \boldsymbol{q}^i_2$. It's natural to insist that the utility of the outcomes has been too simplified by the sums $\boldsymbol{p}^i_j \dotplus \boldsymbol{q}^i_j$. The proper response is that the expected values of the two functions are being compared as a matter of belief, not value or utility. Thus it would seem quite natural to bet that the expected value of $\boldsymbol{p}^i_1 \dotplus \boldsymbol{q}^i_1$ will be greater than that of $\boldsymbol{p}^i_2 \dotplus \boldsymbol{q}^i_2$, no matter how one feels about the relative desirability of sunny versus warm weather. (1982, 433)

And in regards to the non-mixed iterated indicator functions, $n\boldsymbol{p}^i$ where $n > 1$, Suppes offers the following interpretation:

> From an intuitive estimation or gambling standpoint, it's much easier to reflect on the subjective probability of $n\boldsymbol{p}^i$ than of $n\boldsymbol{p}^i + m\boldsymbol{q}^i$. For example, if $\boldsymbol{p}^i(\omega) = 1$ means "heads" in a toss of a coin with unknown bias, then $5\boldsymbol{p}^i$ is just the estimate of 5 such tosses being "heads". (2014, 53)

The intuitive example is selectively chosen. Supposing that Ω is a set of possible worlds, $\boldsymbol{p}^i(\omega) = 1$ in general means that some proposition, p, is true at the world ω. It is not clear how something along these lines will make intuitive sense when p is *there are dogs* or *most roses are red*.

5.3 Multiprimitive structures

Suppose we identify an agent's unconditional probabilities with their probabilities conditional on the logically necessary proposition Ω. Given that, we can see the two approaches just discussed as alterantives ways of enriching the relatively simple systems of comparative (unconditional) confidence $\langle \mathcal{A}, \succsim \rangle$ characterised by Definition 4.1. The first extends the domain of the confidence relation to $\mathcal{A} \times \mathcal{A}$, such that the agent's unconditional confidence relation falls out as a special case; whereas the qualitative expectations approach extends the domain from $\mathcal{A}$ to $\boldsymbol{\mathcal{A}}^i$. The next approach also enriches those systems, though in a different way again: by adding new psychological primitives.

There is an absurd variety of ways this might go, depending on what primitives we choose to add and the structures we take them to possess. One might, for example, add a primitive unary property corresponding to *certainty*. This would be a natural move if one supposes that the property of *being certain that p* might sometimes come apart from the property of *being at least as confident that p as any other proposition*, where the latter is the only feasible way to define the former if one is restricted to definitions in terms of a single confidence relation. Similarly, if one supposes that all-or-nothing belief cannot be characterised in terms of comparative confidence, but still seeks to represent all-or-nothing beliefs within a numerical framework, then one might try adding a primitive *believes* relation by which to do so.

Probably the most commonly suggested additional primitive is an *independence* relation (e.g., Domotor 1970; Fine 1973; Kaplan & Fine 1977; Luce 1978; Luce & Narens 1978; Joyce 2010). Per usual, we say that p and q are independent relative to a probability measure μ whenever

$$\mu(p \cap q) = \mu(p) \cdot \mu(q)$$

In cases where a comparative confidence relation admits of more than one probabilistic representation, which propositions count as probabilistically independent of one another can vary depending on which measures are chosen. An example: suppose $\mathcal{A}$ contains four atoms, a_1–a_4, and the probability measures μ and μ' are defined like so

$$\mu(a_1) = 0.02, \quad \mu(a_2) = 0.08, \quad \mu(a_3) = 0.18, \quad \mu(a_4) = 0.72$$
$$\mu'(a_1) = 0.03, \quad \mu'(a_2) = 0.08, \quad \mu'(a_3) = 0.18, \quad \mu'(a_4) = 0.71$$

The resulting measures are ordinally-equivalent, as represented in the following table:

	μ	μ'		μ	μ'
Ω	1	1	$p_{\langle 123\rangle}$	0.28	0.29
$p_{\langle 234\rangle}$	0.98	0.97	$p_{\langle 23\rangle}$	0.26	0.26
$p_{\langle 134\rangle}$	0.92	0.92	$p_{\langle 13\rangle}$	0.20	0.21
$p_{\langle 34\rangle}$	0.90	0.89	a_3	0.18	0.18
$p_{\langle 124\rangle}$	0.82	0.82	$p_{\langle 12\rangle}$	0.10	0.11
$p_{\langle 24\rangle}$	0.80	0.79	a_2	0.08	0.08
$p_{\langle 14\rangle}$	0.74	0.74	a_1	0.02	0.03
a_4	0.72	0.71	$\varnothing$	0	0

Observe that $p_{\langle 24\rangle}$ and $p_{\langle 34\rangle}$ are independent relative to μ, not relative to μ':

$$\mu(a_4) = \mu(p_{\langle 24\rangle}) \cdot \mu(p_{\langle 34\rangle}), \qquad \mu'(a_4) \neq \mu'(p_{\langle 24\rangle}) \cdot \mu'(p_{\langle 34\rangle})$$

So probabilistic independence is not, in general, *meaningful* relative to the probabilistic measurement of comparative confidence. Since independence is one of the more central concepts in probability theory, and does important theoretical work, we should want to rectify this situation.

One might suppose we could simply solve the problem by imposing further axioms on $\succsim$, thus ensuring a unique probabilistic representation. But this response is inadequate. For one thing, it doesn't solve the problem. Even supposing $\langle \mathcal{A}, \succsim; \sqcup\rangle$ has a unique *probabilistic* representation in $\langle \mathbb{R}, \geq; +\rangle$, there will be many other ways of representing that system in $\langle \mathbb{R}, \geq; +\rangle$—and independence will not be meaningful relative to the broader class of additive homomorphisms where $\varphi(\Omega)$ need not equal 1. (See Luce *et al.* 1990, 277–8, for useful discussion on this point.) Moreover, a system of qualitative probability that's not uniquely probabilistically representable presumably still corresponds to at least one possible system of belief, and the ordinally-equivalent probability measures that represent such systems will play differentiable roles in epistemology and decision theory that we should like to be able to account for.

The better response, then, is to find a system of primitives that will guarantee meaningfulness for independence. One way to do this is to simply add a qualitative independence relation alongside comparative confidence. Let $\perp$ designate a binary relation on $\mathcal{A}$. The goal will be to supply conditions on an enriched system $\langle \mathcal{A}, \succsim, \perp\rangle$ sufficient for the existence of a measure such that

i. $p \succsim q$ iff $\varphi(p) \geq \varphi(q)$
ii. If $p \cap q = \varnothing$, then $\varphi(p \cup q) = \varphi(p) + \varphi(q)$
iii. $p \perp q$ iff $\varphi(p \cap q) = \varphi(p) \cdot \varphi(q)$

Obviously, if $\mathcal{A}$ is finite then a measure will exist only if $\succsim$ satisfies the axioms from Definition 4.1. Necessary axioms for $\perp$ on this interpretation are provided by Suppes (2014):

1. $p \perp \Omega$
2. If $p \perp p$, then $p \sim \Omega$ or $p \sim \varnothing$
3. If $p \perp q$, then $q \perp p$
4. If $p \perp q$, then $p \perp \neg q$
5. If $q \cap r = \varnothing$, and $p \perp q$, $p \perp r$, then $p \perp q \cup r$

Adding such a relation into a system of qualitative probability will in some cases let us meaningfully differentiate between ordinally-equivalent probability measures. It does so for the example just above, for instance, depending on whether $p_{\langle 12 \rangle} \perp p_{\langle 23 \rangle}$. But not always. Consider again the case mentioned in §3.4. We suppose $\mathcal{A} = \{\Omega, p, \neg p, \varnothing\}$, and

$$\Omega \succ p \succ \neg p \succ \varnothing$$

These comparative confidences can be represented by many measures satisfying properties i and ii, provided

$$\varphi(\Omega) > \varphi(p) = \left[\varphi(\Omega) - \varphi(\neg p)\right] > \varphi(\varnothing) = 0$$

Adding any qualitative independence relation and requiring the measure to also satisfy property ii forces those measures to satisfy *normalisation*, and hence forces them to be probability measures. However, it does nothing at all to sort between the many ordinally-equivalent probability measures that fit with those comparative confidences.

It is possible to add yet further primitives that will help to guarantee unique probabilistic representability even where the conditions on $\succsim$ and $\perp$ alone are not enough. For example, Suppes (2014, 49–50) shows that if one adds a primitive *entropic uncertainty* relation $\succsim^u$ (defined over partitions of Ω) and appropriate axioms connecting $\succsim$, $\perp$, and $\succsim^u$, then one can guarantee a unique (absolute scale) representation of the resulting system that happens to be a probability measure. No doubt there are many other primitives one could try including alongside $\succsim$ and $\perp$ that might work too—the matter has, at this stage of the literature, only undergone the most cursory exploration.

Chapter 6

Decision-Theoretic Approaches

A *decision-theoretic representation* is a kind of conjoint representation, typically of a single binary preference relation that decomposes into a representation of beliefs and a representation of desires that pairwise determine those preferences according to a pre-specified decision rule.

Decision-theoretic representations vary along multiple dimensions, depending on the kinds of primitives used to construct the qualitative system, the desired constraints on the numerical representations of belief and desire, or the details of the decision rule. By far the most well-known theorems in this space are those for subjective expected utility theory; here we find the seminal works of Ramsey (1931) and Savage (1954). But there are dozens if not hundreds of variations on these theorems, and many more indeed for the substantial number of alternative non-expected utility theories that have been proposed as descriptive or normative rivals to expected utility theory.

I won't attempt to cover all the variety in this chapter. Instead, I'll start with a brief overview of the main frameworks in which decision-theoretic representations tend to be constructed (§6.1), after which I'll go into more detail on one relatively simple theorem based closely on Ramsey's (§6.2). Finally, I discuss meaningfulness in the conjoint measurement of belief and desire (§6.3), and rebut a few common complaints about the decision-theoretic approach (§6.4).

6.1 The objects of preference

Before we can construct a conjoint representation of preferences as determined by beliefs and desires, we require first an appropriate way of formalising the objects over which the preference relation is to be defined. These objects are variably referred to as *gambles*, *bets*, *prospects*, *options*, *acts*, *decisions*, *choices*, and more, depending on the intended interpretation of the theorem and the personal inclinations of its authors.

Broadly speaking, there are three ways to formalise the objects of preference. These can be ordered by the amount of internal structure they represent those objects as having. At one end of the spectrum are theorems that, like Savage's (1954), employ more or less arbitrary associations between states of nature and consequences. In this context preferences are usually interpreted as being defined over actions the agent might perform, or perhaps intentions to perform those action, with the idea being that actions can be represented by their possible outcomes relative to the states of the world under which the action brings about those outcomes. Where $\mathcal{S} = \{s_1, s_2, \ldots\}$ is a partition of Ω representing different states the world might be in, and $\mathcal{C} = \{c_1, c_2, \ldots\}$ is a set of consequences that any potential action could bring about depending on which state happens to be true, we let each action be represented by a function from $\mathcal{S}$ to $\mathcal{C}$. Then, if f is the function that pairs s_i with c_i, then it represents the action such that were it performed, then if s_1 is the true state then c_1 would result, and if s_2 is the true state then c_2 would result, and so on.

Preferences are then defined over a set of these functions, and a conjoint representation is constructed that typically decomposes into two measures—a function defined over $\mathcal{C}$, and a function defined over an algebra of propositions (usually called 'events') constructed from the states in $\mathcal{S}$. For example, suppose that $\mathcal{S}$ is finite, and $\mathcal{E} = \{e_1, e_2, \ldots\}$ is the algebra of propositions with atoms given by $\mathcal{S}$. Then, an ordinary expected utility theorem provides axioms on a preference relation $\succsim$ defined over a space of actions $\mathcal{C}^{\mathcal{S}}$ sufficient for the existence of a probability measure β on $\mathcal{E}$ ('β' for beliefs) and a real-valued function δ on $\mathcal{C}$ ('δ' for desires) such that $f \succsim g$ if and only if

$$\sum_{s \in \mathcal{S}} \beta(s)\delta[f(s)] \geq \sum_{s \in \mathcal{S}} \beta(s)\delta[g(s)]$$

Note that β and δ must be defined for distinct sets—indeed, in Savage's construction they are disjoint. The reason is that a proposition counts as an event just in case it's logically equivalent to a disjunction of states; hence any proposition that's consistent with any state and its negation cannot be an event. Given that, observe that consequences cannot in general be events if the functional representation of actions is to be at all coherent. We cannot say that f is the action that brings consequence c_i at state s_i, whereas g is the action that brings some other consequence c_j at s_i, if every state s_i logically determines that a particular consequence obtains. For a similar reason, states cannot in general determine actions. Hence, the domain of the belief function cannot include propositions that specify actions under deliberation nor the consequences thereof. For some this is seen as a good-making feature of Savage's construction (e.g., Spohn 1977); for others not so much (e.g., Hájek 2016; Elliott 2017a).

Savage's approach contrasts with the formalisation strategy found at the other end of the spectrum. Here we find theorems that, like Jeffrey's (1965; 1978; see also Bolker 1967 and Domotor 1978), define preferences over an algebra of propositions that simultaneously serves as the domain of both the belief and desire functions. For this reason they are sometimes called 'monoset theorems'. Jeffrey's own theorem supplies axioms on a preference relation $\succsim$ defined over an algebra of proposition $\mathcal{A}$ sufficient for the existence of a probability measure β and a real-valued δ where $p \succsim q$ iff $\delta(p) \geq \delta(q)$, and for $p \in \mathcal{A}$, if $\{p_1, p_2, \ldots, p_n\}$ is any finite partition of p then:

$$\delta(p) = \sum_{i=1}^{n} \beta(p_i \,|\, p)\delta(p_i)$$

It makes little sense to interpret the objects of Jeffrey's preference relation as actions. Some of the propositions in $\mathcal{A}$ may well correspond to actions the agent may choose to perform—these Jeffrey (1968, 170) refers to as *actual* propositions, and decision-making is construed as choosing which of a set of mutually exclusive and jointly exhaustive actual propositions to make true. But many more of the propositions over which $\succsim$ is defined will correspond to no plausible object of choice in any decision context, so $\succsim$ is much better seen in this case as a relative desirability relation:

> To say that p is ranked higher than q [in the agent's preference ordering] means that the agent would welcome the news that p is true more than he would the news that q is true: p would be better news than q. (Jeffrey 1990, 82)

The difference in how the objects of preference are represented is also important from a measurement-theoretic perspective. For any numerical representation of any relation, if that representation is going to be more than just an ordinal scale then one needs appeal to *some* additional formal relation between the relata when characterising the qualitative system—else there will be nothing for the extra-ordinal structure of the representation to be a representation *of*. In the Savage framework, the additional structure is present in the act-functions. For example, Savage's theorem requires:

- If $f(s) = c_1$ and $g(s) = c_2$ for all $s \in \mathcal{S}$, and $f \succ g$, then if $f'(s) = g'(s)$ for all $s \in \mathcal{X} \subset \mathcal{S}$ and otherwise $f'(s) = f(s)$ and $g'(s) = g(s)$, then $f' \succ g'$

In other words, if two acts f' and g' have identical consequences for a subset of states, and for all other states f' has better consequences, then f' should be preferred to g'. In the Jeffrey framework, however, the relata of $\succsim$ have no particular internal structure, they are just sets of worlds. Hence, we need appeal instead to logical (or set-theoretic) relations between propositions. For example, Jeffrey's theorem requires:

- If $(p \cup q) \sim q$ for some $q \in \mathcal{A}$ such that $p \cap q = \varnothing$ and either $p \succ q$ or $q \succ p$, then $(p \cup r) \sim r$ for all $r \in \mathcal{A}$

In other words, if p makes no contribution to the desirability of $p \cup q$ for disjoint p and q of distinct desirabilities, then the agent must presumably have zero confidence in p, and hence for consistency p should make no contribution to the desirability of $p \cup r$ for any other r.

A middle-ground is provided by the third kind of framework, originating with Ramsey (1931), where preferences are defined over a domain of very simple prospects of the form "c_1 if p, c_2 otherwise". These are typically interpreted as conjunctions of conditionals, perhaps corresponding to potential choices or gambles the agent might take, and formalised as n-tuples of conditions and consequences—e.g., (c_1, p, c_2). Most such theorems focus on binary prospects like the one just described. In some cases preferences are also defined for ternary prospects "c_1 if p, c_2 if q, c_3 otherwise", or even quarternary prospects, but nothing nearly so richly structured as the (potentially infinitary) act-functions we find in Savagean frameworks. (For examples of theorems in the Ramseyan framework, see Debreu 1959; Davidson & Suppes 1956; Davidson *et al.* 1957; Fishburn 1967; Elliott 2017b; 2017c.) The theorem discussed in the next section belongs to this third class; it is a modified version of Ramsey's original.

Before we move on, it's worth noting that as a framework for formalising a decision theory, the Ramseyan approach is *extremely* limited—especially in contrast to either of the Savagean or Jeffreyan frameworks just discussed. Most decision contexts involve choices between options that cannot plausibly be reduced to simple n-ary prospects for very small n. Consequently, if the goal of the representation theorem were to provide a complete and fully general axiomatisation of a decision theory in terms of preferences, then the Ramseyan framework would be grossly inadequate. This is what Savage and Jeffrey sought to achieve with their works, but for a theory of measurement we needn't ask so much. The purpose of the construction is rather to isolate a qualitative conjoint psychological system with a relational structure that suffices to explain the quantitation of belief.

With that in mind, there need be no assumption that the system in question should include the entirety of the agent's preferences nor that the decision rule should be generally applicable to all conceivable decision situations. It may be, of course, that the quantitation of belief could *also* be explained in the context of some richer preference structure given a much more general decision rule capable of handling more varied prospects than simple binary prospects. The claim here is just that we don't *need* to consider a full suite of preferences to explain why it makes sense to measure beliefs the way we do.

6.2 Ramsey's theorem

The picture that I'll be painting here will be very similar to the one that was sketched in §3.2. From a single preference ordering over a space of simple binary prospects of the form "c_1 if p, c_2 otherwise", the goal is to extract numerical representations of belief and desire that conjointly represent those preferences.

The first step is to be precise about the form of the representation. We let $\succsim$ be a preference relation defined over a set $\mathcal{G}$ of prospects. Where $\mathcal{A} = \{p, q, r, \dots\}$ is an algebra of propositions and $\mathcal{C} = \{c_1, c_2, c_3, \dots\}$ is a set of consequences, we formalise prospects as 3-tuples (c_1, p, c_2) in $\mathcal{G} \subseteq \mathcal{C} \times \mathcal{A} \times \mathcal{C}$. We needn't assume that $\mathcal{A}$ and $\mathcal{C}$ are disjoint sets, nor do we have to assume that the consequences are maximally specific in any particular sense.[1] For simplicity, though, we will assume that both $\mathcal{A}$ and $\mathcal{C}$ are finite. This will let us ignore a complicated 'Archimedean' axiom that's trivially satisfied in finite contexts.

We desire a function $\varphi : \mathcal{G} \mapsto \mathbb{R}$ that represents $\succsim$ in the sense that

$$(c_1, p, c_2) \succsim (c_3, q, c_4) \text{ iff } \varphi\big[(c_1, p, c_2)\big] \geq \varphi\big[(c_3, q, c_4)\big],$$

[1] In Ramsey's essay, consequences are ambiguously said to be maximally specific worlds, plus in some cases *almost*-worlds that are maximally specific up to a single question about which the agent cares not. With some minor adjustments this ends up being unnecessary for the representation result and for the decision theory underlying it.

where φ itself decomposes into two functions $\beta : \mathcal{A} \mapsto \mathbb{R}$ ('β' for *beliefs*) and $\delta : \mathcal{C} \mapsto \mathbb{R}$ ('δ' for *desires*) such that

$$\varphi[(c_1, p, c_2)] = \delta(c_1)\beta(p) + \delta(c_2)[1 - \beta(p)]$$

Call this the *simplified formula.*

Note an immediate complication: the simplified formula is too simple! It implies that the three factors contributing to the value of a prospect are independent of one another. In particular, according to the simplified formula, $\beta(p) = \beta(q)$ implies $\varphi[(c_1, p, c_2)] = \varphi[(c_1, q, c_2)]$. However, the desirability of c_1 may vary considerably depending on whether it obtains in a context where p is true versus a context where q is true, say, and likewise for c_1 and $\neg p$ versus $\neg q$, leading to a difference in desirability between (c_1, p, c_2) and (c_1, q, c_2). In general, the value of a prospect's consequences ought to be judged relative to the conditions under which those consequences obtain. Consequently, the value of (c_1, p, c_2) isn't given by the simplified formula, but should instead be given by a slightly more complicated:

$$\varphi[(c_1, p, c_2)] = \delta(p \wedge c_1)\beta(p) + \delta(\neg p \wedge c_2)[1 - \beta(p)]$$

The implication is that if $\succsim$ is defined for all possible prospects in $\mathcal{C} \times \mathcal{A} \times \mathcal{C}$, and is determined by this more complicated decision rule, then it *cannot* be represented in the desired manner. And the resolution is to restrict $\mathcal{G}$ to those prospects (c_1, p, c_2) such that the agent is indifferent between c_1 and $c_1 \wedge p$, and likewise between c_2 and $c_2 \wedge \neg p$, in which case the complicated formula reduces to the simplified formula. If propositions are coarsely-individuated—if they are sets of logically possible worlds, say—then a natural way to achieve this restriction is to suppose that a prospect is in $\mathcal{G}$ only if its consequences entail the conditions under which they obtain; in this case, $c_1 = c_1 \wedge p$ and $c_2 = c_2 \wedge \neg p$. More precisely:

1. $(c_1, p, c_2) \in \mathbf{G}$ only if c_1 implies $p \wedge c_1$ and c_2 implies $\neg p \wedge c_2$; and $p \wedge c_1$ is inconsistent only if p is inconsistent, and $\neg p \wedge c_2$ is inconsistent only if $\neg p$ is inconsistent (*restricted prospects*)

We will also need to ensure that $\mathcal{G}$ is rich enough to ensure the existence of the desired representation. There are five richness axioms in total, starting with:

2. For every $c \in \mathcal{C}$, there is a prospect $(c, p, c) \in \mathcal{G}$ (*trivial prospects*)

The sole purpose of this axiom is to let us extend the preference ordering $\succsim$ to consequences, in the obvious way:

$$c_1 \succsim c_2 \text{ iff } \exists (c_1, p, c_1), (c_2, q, c_2) \in \mathcal{G} : (c_1, p, c_1) \succsim (c_2, q, c_2)$$

Axioms 7–10, below, will ensure that $\succsim$ on $\mathcal{C}$ is a weak order.[2]

[2] The *trivial prospects* axiom is isn't necessary if we treat preferences over consequences as a primitive. This is what Ramsey did. However, letting preferences be defined in the first instance over only $\mathcal{G}$, rather than both $\mathcal{C}$ and $\mathcal{G}$, makes some parts of the construction slightly more natural.

Before I state the four remaining richness axioms, some notation will prove useful. First, let $\underline{c}$ designate the set of consequences c' in $\mathcal{C}$ such that $c \sim c'$. In terms of the intended representation, $\underline{c}$ contains all and only those c' such that $\delta(c) = \delta(c')$. We then use $(\underline{c_1}, p, \underline{c_2})$ for an arbitrary prospect conditional on p with consequences equal in desirability to c_1 and c_2. Next, suppose that for $c_1 \succ c_2$,

$$(\underline{c_1}, p, \underline{c_2}) \succsim (\underline{c_1}, q, \underline{c_2})$$

Supposing that $\succsim$ is represented in the intended fashion, it must follow that $\beta(p) \geq \beta(q)$. Consequently, if

$$(\underline{c_1}, p, \underline{c_2}) \succsim (\underline{c_1}, \neg p, \underline{c_2})$$

then $\beta(p) = \beta(\neg p)$. In this fashion we can isolate the half-probability propositions in $\mathcal{A}$. We use $(\underline{c_1}, \frac{1}{2}, \underline{c_2})$ for a prospect with consequences equal in desirability to c_1 and c_2 conditional on one or another of these half-probability propositions. These prospects are used to define halfway points between the desirabilities of c_1 and c_2. Finally, we characterise a qualitative ordering $\succsim_\Delta$ over $\mathcal{C} \times \mathcal{C}$ like so:

$$(c_1, c_2) \succsim_\Delta (c_3, c_4) \text{ iff } (\underline{c_1}, \frac{1}{2}, \underline{c_4}) \succsim (\underline{c_2}, \frac{1}{2}, \underline{c_3})$$

Defined as such, $\succsim_\Delta$ represents the relative size of intervals in desirability—in terms of the final representation, we will see

$$(c_1, c_2) \succsim_\Delta (c_3, c_4) \text{ iff } \delta(c_1) - \delta(c_2) \geq \delta(c_3) - \delta(c_4)$$

We can now state the remaining richness axioms with ease:

3. If $c_1 \succ c_2$ then there is some $(\underline{c_1}, \frac{1}{2}, \underline{c_2}) \in \mathcal{G}$ (*halfway prospects*)
4. If $(c_1, c_2) \succsim_\Delta (c_3, c_4) \succsim_\Delta (c_1, c_1)$, then there are $c_5, c_6 \in \mathcal{C}$ such that $(c_1, c_5) \sim_\Delta (c_3, c_4) \sim_\Delta (c_6, c_2)$ (Δ-*solvability*)
5. For every $(c_1, p, c_2) \in \mathcal{G}$, there's some $(c_3, q, c_3) \in \mathcal{G}$ such that $(c_1, p, c_2) \sim (c_3, q, c_3)$, or some $(c_4, \frac{1}{2}, c_5) \in \mathcal{G}$ such that $(c_1, p, c_2) \sim (\underline{c_4}, \frac{1}{2}, \underline{c_5})$ (*extendibility*)
6. For every $p \in \mathcal{A}$, there's some $(\underline{c_1}, p, \underline{c_2}) \in \mathcal{G}$ such that $c_1 \succ c_2$ or $c_2 \succ c_1$ (*non-trivial prospects*)

The *halfway prospects* axiom ensures that we can always define halfway points between the desirabilities of any two consequences; this is what lets us characterise the $\succsim_\Delta$ ordering of interval sizes in desirability. Δ-*solvability* is a non-necessary condition used to guarantee that for any non-zero interval in desirability between two consequences there's another interval of the same size to which it can be 'added'; this allows ratios of differences to be readily defined, which is what ultimately allows the construction of an interval-scale measure δ. The *extendibility* axiom is then used to extend δ on $\mathcal{C}$ to all of $\mathcal{G}$, and hence define φ. Finally, *non-trivial prospects* is used to ensure that there are enough prospects around such that a degree of belief can be defined for every proposition in $\mathcal{A}$.

Definition 6.1 Where $\mathcal{A}$ is a finite algebra of propositions and $\mathcal{C}$ is a finite set of propositions, $\langle \mathcal{G}, \succsim \rangle$ is a *finite Ramseyan structure* iff $\mathcal{G} \subseteq \mathcal{C} \times \mathcal{A} \times \mathcal{C}$ satisfying *restricted prospects*, *trivial prospects*, *halfway prospects*, Δ-*solvability*, *extendability* and *non-trivial prospects*.

What remains is to specify conditions on a finite Ramsey structure sufficient for the existence of the desired representation. We proceed in three stages, starting with the construction of the desirability function δ, which uses:

7. $\succsim$ is a weak order (*weak order*)
8. $\succsim_\Delta$ is transitive (Δ-*transitivity*)
9. For all $c_1, c_2 \in \mathcal{C}$, $(\underline{c_1}, \frac{1}{2}, \underline{c_2}) \sim (\underline{c_2}, \frac{1}{2}, \underline{c_1})$ (*reversibility*)
10. For all $c_1, c_2 \in \mathcal{C}$, if $c_1 \succsim c_2$, then $(\underline{c_1}, p, \underline{c_1}) \succsim (\underline{c_1}, q, \underline{c_2})$ (*averaging*)

In overview, δ is derived as follows. (For details, see Elliott 2017c and Krantz *et al.* 1971, 145–52.) First, we use *weak order*, Δ-*transitivity* and *reversibility*, in conjunction with *halfway prospects* and Δ-*solvability*, to define a concatenation operation $\oplus$ over intervals of desirability such that $\langle \mathcal{C} \times \mathcal{C}, \succsim_\Delta ; \oplus \rangle$ is an additive extensive structure, allowing for a ratio-scale measure of desirability intervals. Given *averaging*, we can then define an interval-scale measure $\delta : \mathcal{C} \mapsto \mathbb{R}$ such that:

i. $c_1 \succsim c_2$ iff $\delta(c_1) \geq \delta(c_2)$
ii. $(c_1, c_2) \succsim_\Delta (c_3, c_4)$ iff $\delta(c_1) - \delta(c_2) \geq \delta(c_3) - \delta(c_4)$

Next, we characterise φ over $\mathcal{G}$ using δ over $\mathcal{C}$, like so:

$$\varphi[(c_1, p, c_2)] = \begin{cases} \delta(c_3) & \text{if } (c_1, p, c_2) \sim (\underline{c_3}, q, \underline{c_3}) \\ \frac{1}{2}[\delta(c_3) - \delta(c_4)] & \text{if } c_3 \succ_\delta c_4 \text{ and } (c_1, p, c_2) \sim (\underline{c_3}, \frac{1}{2}, \underline{c_4}) \end{cases}$$

The *extendibility* axiom ensures that the definition is adequate. Note, of course, that $\delta(c) = \varphi[(c, p, c)]$. Finally, we extract the belief function β directly out of φ. Where $c_1 \succ c_2$, reorganising the simplified formula above gets us a definition of β as a ratio of differences in desirability:

$$\beta(p) = \frac{\varphi[(c_1, p, c_2)] - \delta(c_2)}{\delta(c_1) - \delta(c_2)}$$

This last step requires *non-trivial prospects*, plus one final axiom that essentially asserts that the contribution $\beta(p)$ makes to the overall value of a prospect is independent of the desirabilities of its consequences. The upshot will be that the above definition of $\beta(p)$ won't depend on the choice of prospect on p.

Expressed directly in terms of preferences, this axiom is quite complicated and not at all intuitive—the interested reader should refer to Davidson & Suppes' (1956) axiom A10 and the associated definitions for how it goes. We can simplify matters greatly by expressing it instead in terms of the intended representation:

11. For all $p \in \mathcal{A}$, if $\delta(c_1) \neq \delta(c_2)$, $\delta(c_3) \neq \delta(c_4)$, and $(c_1, p, c_2), (c_3, p, c_4) \in \mathcal{G}$, then
$$\frac{\varphi[(c_1, p, c_2)] - \delta(c_2)}{\delta(c_1) - \delta(c_2)} = \frac{\varphi[(c_3, p, c_4)] - \delta(c_4)}{\delta(c_3) - \delta(c_4)}$$
(*independence*)

Putting that all together:

Theorem 6.1 (Ramsey 1931; Elliott 2017c) *Where* $\mathcal{G} \subseteq \mathcal{C}\times\mathcal{A}\times\mathcal{C}$, *suppose that* $\langle \mathcal{G}, \succsim \rangle$ *is a finite Ramseyan structure satisfying* weak ordering, Δ-transitivity, reversibility, averaging, *and* independence. *Then there are functions* $\varphi : \mathcal{G} \mapsto \mathbb{R}$, $\delta : \mathcal{C} \mapsto \mathbb{R}$ *and* $\beta : \mathcal{A} \mapsto \mathbb{R}$, *such that*

i. $(c_1, p, c_2) \succsim (c_3, q, c_4)$ *iff* $\varphi[(c_1, p, c_2)] \geq \varphi[(c_3, q, c_4)]$
ii. $\varphi[(c_1, p, c_2)] = \delta(c_1)\beta(p) + \delta(c_2)[1 - \beta(p)]$

Furthermore, δ *is unique up to a positive affine transformation, while* β *is unique and has the property that for all* $p \in \mathcal{A}$,

iii. $1 \geq \beta(p) = 1 - \beta(\neg p) \geq 0$

6.3 Uniqueness and meaningfulness

Note that the β specified in Theorem 6.1 is unique *simpliciter*—an absolute scale. The uniqueness clause applies to all conjoint representations satisfying properties i and ii; property iii, by contrast, is not an explicit stipulation on the form of the representation but is rather derived from that representation. This is a result of how it was defined—as a dimensionless ratio of differences in desirability—and the fact that δ is measured on an interval scale.

It is an often remarked-upon fact that expected utility (EU) representations of a preference ordering are not unique. It is a consequence of Theorem 6.1, for example, that if $\succsim$ has an EU representation involving β and δ, then it will have another such representation involving β and $\delta^\star$, where

$$\delta^\star(c) = 9\delta(c) + 1$$

Since desirabilities are ordinarily understood to be represented on nothing stronger than an interval scale, like temperatures as measured in Celsius or Fahrenheit or years as measured on different calendars, the usual response to this fact is that there is no meaningful difference between δ and $\delta^\star$.

So far so good, but consider next a variation on that same idea (originating with Zynda 2000). If $\succsim$ has an EU representation involving β and δ, then it will also have a representation involving a different (and non-probabilistic) function $\beta^\star$ and δ, where

$$\beta^\star(p) = 9\beta(p) + 1$$

In this case, though, the joint representation employs a slightly different decision rule: $\beta^\star$ and δ are combined according to the *valuation-maximisation rule*. Where γ is an n-ary prospect with consequences $\gamma(p_i)$ conditional on which element of a partition $p_1, \ldots, p_n$ happens to be true, this alternative rule tells us that $\gamma \succsim \gamma'$ just in case

$$\sum_{i=1}^{n} \beta^\star(p_i)\delta(\gamma(p_i)) - \delta(\gamma(p_i)) \;\geq\; \sum_{i=1}^{n} \beta^\star(p_i)\delta(\gamma'(p_i)) - \delta(\gamma'(p_i))$$

The proof that $\succsim$ has an EU representation with β and δ just in case it has a valuation-maximisation representation with $\beta^\star$ and δ is straightforward, and worth noting for discussion later. The key step is just to note that the transformation from β to $\beta^\star$ is invertible; hence:

$$\beta(p) = \frac{\beta^\star(p) - 1}{9}$$

Given that, we know that if $\succsim$ has an EU representation when $\gamma \succsim \gamma'$ iff

$$\sum_{i=1}^{n} \beta(p_i)\delta\big(\gamma(p_i)\big) \;\geq\; \sum_{i=1}^{n} \beta(p_i)\delta\big(\gamma'(p_i)\big)$$

Substituting in the above:

$$\sum_{i=1}^{n} \left(\frac{\beta^\star(p_i) - 1}{9}\right)\delta\big(\gamma(p_i)\big) \;\geq\; \sum_{i=1}^{n} \left(\frac{\beta^\star(p_i) - 1}{9}\right)\delta\big(\gamma'(p_i)\big)$$

Then dropping the constant factor:

$$\sum_{i=1}^{n} \big(\beta^\star(p_i) - 1\big)\delta\big(\gamma(p_i)\big) \;\geq\; \sum_{i=1}^{n} \big(\beta^\star(p_i) - 1\big)\delta\big(\gamma'(p_i)\big)$$

Finally, by multiplying the brackets out we get the valuation-maximisation rule exactly as above.

By analogy with δ and $\delta^\star$, one might imagine that this tells us something important about meaningfulness in β and $\beta^\star$—namely, that just as we should want to say that what's *meaningful* in δ and $\delta^\star$ is what's invariant between them, so too should we want to say that what's *meaningful* in β and $\beta^\star$ is what's invariant between them. As Lyle Zynda has suggested,

> One might point out that [$\beta^\star$] is simply a linear transformation of [β], 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* [β and $\beta^\star$]… Since [$\beta^\star$] is subadditive rather than additive, this would commit [one] to the view that additivity cannot be taken *literally* as a property common to all rational degrees of belief… (2000, 64)

And a little further on, Zynda argues that β and $\beta^\star$ will share a common ordering, and thus represent the same comparative confidences. Hence,

> 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, 65, emphasis added)

But there's two errors here; I'll call them the *shallow error* and the *deeper error*. The example *does* teach us something important about meaningfulness in β, but this is not it.

First the shallow error: while it's true that orders are preserved across β and $\beta^\star$, that's not *all* that's preserved. The positive affine transformation that relates β and $\beta^\star$ also preserves ratios of intervals, and those ratios have a decision-theoretic role to play according to both the expected utility rule and the valuation-maximisation rule. Again, the example in §3.4 suffices to make this point. Where $\mathcal{A} = \{\Omega, p, \neg p, \varnothing\}$, we imagine that Ramsey has a choice between the following bets:

α: receive \$1 if p is true, nothing otherwise
β: receive \$2 if p is false, nothing otherwise

According to expected utility theory, Ramsey ought to weakly prefer α just in case

$$\frac{\beta(\Omega)}{\beta(p)} \leq \frac{\beta(p)}{\beta(\neg p)}$$

Since ordinally-equivalent probability measures differ on whether that inequality holds, expected utility theory draws distinctions between such measures. The same is true for the valuation-maximisation rule, according to which Ramsey should weakly prefer α iff

$$\frac{\beta^\star(\Omega)}{\beta^\star(p)} \leq \frac{\beta^\star(p)}{\beta^\star(\neg p)}$$

The point generalises to any alternate decision-theoretic representation that's 'equivalent' to the EU representation in the same way—whatever the relation between them, more than just ordinal information must be shared.

Do not be tempted, though, to infer from this fact that ratios of intervals are meaningful in β and $\beta^\star$. That would be just as fallacious. For consider yet another alternate decision-theoretic representation. This time, we define $\beta^\dagger$ such that

$$\beta^\dagger(p) = \beta^\star(p)^2 = 81\beta(p)^2 + 18\beta(p) + 1$$

We can then show that $\succsim$ has an EU representation involving β and δ iff it has a *schmaluation-maximisation* representation involving $\beta^\dagger$ and δ, where $\gamma \succsim \gamma'$ iff

$$\sum_{i=1}^{n} \Big(\sqrt{\beta^\dagger(p_i)} - 1\Big)\delta(\gamma(p_i)) \;\geq\; \sum_{i=1}^{n} \Big(\sqrt{\beta^\dagger(p_i)} - 1\Big)\delta(\gamma'(p_i))$$

The proof is nearly identical to that for the equivalence of EU representations and valuation-maximisation representations, but for the last step where instead of multiplying out we just substitute $\beta^\star(p_i)$ for the square root of $\beta^\dagger(p_i)$. Again, since the transformation from $\beta^\star$ to $\beta^\dagger$ (or from β to $\beta^\dagger$) is invertible, one just needs to plug the conversion directly into the decision rule.

Ratios of differences are *not* shared between β and $\beta^\dagger$. So if the inference is that anything not shared across 'equivalent' representations is meaningless, we'd be forced to conclude that difference ratios are meaningless in β. Worse still, we can construct an 'equivalent' representation for *any* β^x related to β by any invertible transformation whatsoever. This includes transformations that do not preserve ratios, or ratios of differences, or even orderings. There is, as such, virtually *nothing* that's shared across all such representations.

You can hopefully now see the deeper error, if you haven't spotted it already: the alternate representations are all 'equivalent' in the sense of being equally legitimate ways to represent a system of preferences, but they are representations in distinct numerical systems—and one cannot infer meaninglessness from any variance between representations across numerical systems. For any additive measure φ of length, for instance, and any invertible transformation τ, there is an 'equivalent' measure φ' in a different numerical system where φ and φ' are related by τ. It would be a gross error to infer from this fact that ratios, difference ratios, and even orderings of length are *meaningless*. We don't make that error in the case of length; we shouldn't make it for beliefs.

What's common to β, $\beta^\star$ and $\beta^\dagger$ are the similar roles they play in the models of decision-making that make use of them respectively—*viz.*, *these* beliefs interact with *those* desires to produce such-and-such preferences. Of course there are many ways to represent that connection. Consider, for an analogy, the relationship between force (F), mass (m), and acceleration (a). Where those are represented in kg·m/s^2, kg, and m/s^2 respectively, that connection is represented $F = ma$. But if we alter the numerical representation of one while holding the others fixed we end up with different 'equivalent' ways of representing that connection. Where mass is measured in grams, $F = \frac{ma}{1000}$. Where mass is measured in pounds, $F = \frac{5ma}{11}$. And where mass is measured on the multiplicative version of the kilogram scale with 1kg$= 2x$, $F = \log_2(x)a$. In all cases the superficial form of the rule has changed, but there's no fundamental difference in the qualitative relation being represented.

The real lesson of the example is that the structure being represented by β in the conjoint representation specified by Theorem 6.1 isn't anything internal to the system of beliefs itself, considered in isolation from anything else, but the relationship that holds between beliefs, desires, and preferences. That is what's invariant, and that is why we cannot transform the belief function without making adjustments to the decision rule—because the *meaning* of β is tied up with how it interacts with δ to produce preferences.

6.4 An apology for representation theorems

Theorem 6.1 describes an extremely flexible representation of belief—β must be such that $\beta(p) \in [0, 1]$ and $\beta(p) = 1 - \beta(\neg p)$, but otherwise there are few constraints on the shape it must take. It's straightforward to construct finite Ramseyan structures such that β is logically non-omniscient, and it's equally straightforward to construct finite Ramseyan structures such that β is a probability measure or otherwise logically omniscient. Given the desiderata discussed in §3.4, I take it that this flexibility is a desirable feature. It allows for a non-disjunctive theory of belief measurement that's consistent with a range of probabilistic and non-probabilistic representations, on a more-than-merely ordinal scale without requiring logical omniscience.

The reason for β's flexibility is that the degree of belief assigned to p is determined independently of almost any other proposition besides $\neg p$. This contrasts with epistemic approaches (and Jeffrey-style decision-theoretic approaches). The quantitation of belief on this picture requires no particular appeal to relations between belief states or the contents thereof, but instead depends on systematic relationships between the agent's degree of belief in p and the value they attach to prospects conditional on p. As Ramsey (1931, 169) put it,

> [The] degree of a belief is a causal property of it, which we can express vaguely as the extent to which we are prepared to act on it.

A rough way to express the difference: on the epistemic approach, the strength of Sally's belief towards p is twice that of q when p is equiprobable with the disjunction of two mutually incompatible q' and q'' equiprobable with q; on the Ramseyan approach, if p is believed to twice the degree as q, this will be manifest in the difference in desirability between (c_1, p, c_2) and c_2 being twice the difference between (c_1, q, c_2) and c_2 (where $c_1 \succ c_2$).

It's worth emphasising again that the explanatory relation between belief and preference needn't be a *constitutive* relation. Many have claimed to find in Ramsey's essay the thesis that beliefs are nothing over and above preferences as manifest in choice dispositions, but Ramsey himself characterised the relationship between them in causal terms. More generally, what's required is that the numerical representation has a qualitative interpretation in terms of *some* systematic relation between beliefs and preferences, and that relation may take many forms. Nothing about Theorem 6.1 should be taken to imply that beliefs are *reducible* to preferences. Yes, when proving the theorem we first characterise a desirability function that represents preferences and from that derive a belief function—but it's absurd to infer any kind of ontological or conceptual ordering between quantities just from the order in which their numerical representations happen to be derived in a conjoint representation thereof.

Recognising this fallacy helps in dealing with some of the common objections to the decision-theoretic approach. An exemplar here is Eriksson & Hájek's (2007) Zen monk example. A 'Zen monk' is an agent who is indifferent between all consequences, and therefore indifferent between all prospects. The preferences of such an agent would violate *non-trivial prospects* in such a way that β would be undefinable for all propositions. Yet, presumably, such an agent could still have determinate degrees of belief, and two Zen monks could have distinct degrees of belief between them. If such beings could exist, then they are a counterexample to the thesis that an agent's degrees of belief are nothing over and above their preferences. But the zen monk is much less problematic if we think that the strength of an agent's belief as "a causal property of it" that need not be manifest in all cases (Elliott 2019a). Even if a zen monk is indifferent among all consequences, she may still be in a state of belief the typical causal role of which would only become apparent if she were no longer universally indifferent. What it is to believe p to degree x, on this picture, is to be in a state whose typical causal role in connection to preferences and desire is reflected in the class of systems with representations such that $\beta(p) = x$.

Another common concern is that a theorem like Ramsey's only establishes conditions under which a preference relation behaves *as if* it's determined by such-and-such beliefs and desires combined according to the expected utility rule—it doesn't guarantee that the preferences *really are* the product of those beliefs and desires (cf. Zynda 2000; Christensen 2001; Eriksson & Hájek 2007; Meacham & Weisberg 2011). The observation is correct, of course, just as it's true that a theorem for the conjoint measurement of momentum as determined by mass and velocity only supplies conditions under which momentum behaves *as if* it's determined by mass and velocity. So what? If the point of the representation theorem were to show that an agent whose preferences satisfy the stated axioms thereby has the beliefs and desires they are thereby represented as having as a matter of metaphysical or conceptual necessity, then it would be safe to say that they do no such thing. Lucky, then, that this not the only way to interpret decision-theoretic representation theorems!

A much more fruitful interpretation is in terms of measurement. The aptness of the representation is *presupposed*, not derived from the representation theorem; what the latter supplies is an explanation of the quantitation of belief and desire in the context of that particular model of decision-making:

> I propose to take as a basis a general psychological theory, which … comes, I think, fairly close to the truth in the sorts of cases with which we are most concerned. I mean the theory that we act in the way we think most likely to realize the objects of our desires, so that a person's actions are completely determined by his desires and opinions. (Ramsey 1931, 173)

There's nothing special about the decision-theoretic approach in this respect—the quantitation of any quantity by any method is always explained against a backdrop of theoretical presuppositions.

A natural followup concern, though, is that the presupposed psychological model itself is unrealistic. If that's the case, then the qualitative systems these models represent will presumably fail to capture any explanatorily relevant relations at all. Addressing this concern takes a little more work. To start, let me flag that there will be two main sources for any apparent lack-of-realism. The first is the very precise nature of the intended numerical representation—it requires real-valued degrees of belief and desire, combined with perfect consistency according to a precisely specified decision rule, and to achieve this we usually need strong assumptions about the richness of the domain over which preferences are defined and the precise structure of the preferences over that domain. The second source will be the structure of the decision rule itself—perhaps, for instance, we do not simply evaluate prospects by weighing the values of its consequences against our confidence those consequences will obtain, but instead also take risk into account in a manner that cannot properly be captured by the expected utility rule (or any 'equivalent' rule).

I'm disinclined to worry much about the lack of realism arising due to the former source. It's true that very strong axioms are usually required to establish the uniqueness results these theorems are famous for. That's hardly surprising—

uniqueness is an extraordinarily strong property for a real-valued representation of belief to have! Such is an inevitable consequence of trying to model a squishy psychological system in a rigid numerical framework, and any feasible theory of belief measurement needs allow for some idealisations that make the topic tractable. It's enough if the systems we characterise are in the ballpark of realism. More importantly, it's usually possible without great difficulty to isolate and weaken or remove those axioms (or parts of the axioms) that are involved in fixing the precision of the representation, especially if we're willing to accept somewhat weaker uniqueness conditions as a result.

I said 'somewhat weaker' for a reason. Critics of decision-theoretic representation theorems tend to argue as if failing to guarantee a unique real-valued belief function is equivalent to establishing no bounds on degrees of belief at all—as if any lack-of-uniqueness implies radical non-uniqueness. And that just isn't the case. Dropping those axioms need not, and does not in general imply radical indeterminacy. Consider some examples. I've already (in §3.3) talked about how the presumed completeness of $\succsim$, implied by the *weak order* axiom, can be dropped so as to allow for a representation of 'imprecise' beliefs and desires. So consider instead the *extendibility* axiom. This (structural) axiom helps us to pinpoint precise degrees of belief by fixing a precise φ-value for some appropriate prospect (c_1, p, c_2) conditional on p; it does so either by setting that value equal to the desirability of an consequence or equal to the midway point between two consequences. But where *extendibility* is violated and the required prospects don't exist, we can still characterise bounds on degrees of belief provided there are c_3, c_4 such that

$$c_1 \succsim c_3 \succ (c_1, p, c_2) \succ c_4 \succsim c_2$$

In this case, the value of $\beta(p)$ will be bound like so:

$$\frac{\delta(c_1) - \delta(c_3)}{\delta(c_1) - \delta(c_2)} > \beta(p) > \frac{\delta(c_1) - \delta(c_4)}{\delta(c_1) - \delta(c_2)}$$

More or less the same effect can be achieved if the *independence* axiom is violated. That axiom forces a very precise consistency across how every prospect conditional on p is evaluated relative to its consequences, which is necessary if $\beta(p)$ is to be defined as a real value in the stipulated manner. But it's possible to weaken that axiom to allow for a bit of fuzziness in the evaluation of prospects and corresponding fuzziness in the characterisation of β. Essentially, where the axiom is violated then for every p there's still a unique—and potentially very narrow—interval $[x, y]$ such that every prospect on p is valued as if $\beta(p) \in [x, y]$. And in a similar fashion again we can characterise bounds on the desirabilities of consequences in systems where *trivial gambles*, *halfway prospects* and/or Δ-*solvability* are violated, with consequences for the precision of the belief function defined in terms of those desirabilities.

In general, some of the axioms (or some parts of some axioms) primarily serve to ensure a *precise* numerical representation. They tend to be quite unrealistic, but they can also be weakened or dropped. The effect of doing so is a little less precision in the numbers obtained, but nothing more substantially affecting the basic explanatory structure being represented. I expect that Ramsey understood this point well, and was expressing as much when he wrote:

> I have not worked out the mathematical logic of this in detail, because this would, I think, be rather like working out to seven places of decimals a result only valid to two. (1931, 180)

That, I think, is the right attitude. It's unrealistic to suppose that degrees of belief (and desire) have the all the precision of the real numbers, but we gain some insight into their quantitation by pretending otherwise and lose nothing of great import in the fiction.

It's of more concern if the decision rule itself is psychologically unrealistic, even after accounting for some imprecision in their degrees of belief and desirabilities. We must be a little careful here, though. Suppose that ordinary decision-makers systematically violate the expected utility rule when evaluating prospects. In this case the expected utility rule can still serve as a rational ideal, and Theorem 6.1 may still prove useful in explaining the quantitation of belief by reference to the role that one's beliefs regarding p *ought* to play in connection to how they *ought* to evaluate prospects conditional on p. I said above that we don't have to interpret the systematic relationship between belief and preference that explains the conjoint quantitation thereof as a *constitutive* relation; we don't have to interpret it as a *descriptive* causal relation either. Similarly, an analytic functionalist might say that the expected utility rule captures the essence of folk psychology (*a la* Lewis 1974), and hence a theorem like Ramsey's can help explain how beliefs are quantitated *according to folk psychology*. Since it's no commitment of analytic functionalism that folk psychology provides the optimal (descriptive or normative) account of decision-making, concerns about the adequacy of expected utility theory are largely irrelevant to this interpretation. The theory is uncontroversially *close* to the truth in either case, and the analytic functionalist needs nothing stronger than this.

Still, one may be concerned that the expected utility rule is neither descriptively nor normatively adequate, and may not be satisfied with the analytic functionalist's interpretation. The good news is that there are essentially similar theorems for any number of alternative non-expected utility theories. Alternative representations are far too numerous to discuss in detail, but it's worth looking at one example—Kahneman & Tversky's (1979) prospect theory.[3]

I'll start by describing the theory. We designate a special (non-)consequence the *status quo*; in the representation, the desirability of the status quo will be fixed at zero, hence we'll label it 0. We then focus in on ternary prospects of the form "c_1 if p, c_2 if q, and 0 otherwise", where p and q are mutually exclusive. We assume that degrees of belief are values between zero and one that sum to one for sets of mutually exclusive and jointly exhaustive propositions. Fixing the desirability of the status quo at zero, according to the expected utility rule:

$$\varphi\big[(c_1, p, c_2, q, 0)\big] = \beta(p)\delta(c_1) + \beta(q)\delta(c_2)$$

[3] I note this example because it's simple, because prospect theory is particularly well-known as a descriptive decision theory, and because it's formally similar to expected utility theory's main contemporary normative contender: risk-weighted utility theory (Buchak 2013).

In other words, the part of the prospect corresponding to the status quo makes no contribution to the value of the prospect, which is a weighted average of the desirabilities of the other consequences. According to prospect theory, however, the weights aren't given by the agent's degrees of belief directly. Instead they're given by a *decision weight* that corresponds to the agent's beliefs plus their attitudes towards risk, where the latter modify the impact the agent's degrees of belief have on the overall value of a gamble. Where $\pi : [0,1] \mapsto [0,1]$ and $\pi(0) = 0$ and $\pi(1) = 1$,

$$\varphi[(c_1, p, c_2, 0)] = \pi[\beta(p)]\delta(c_1) + \pi[\beta(q)]\delta(c_2)$$

For example, suppose $\beta(p) = \beta(q) = \frac{1}{2}$, $\delta(c_1) > \delta(c_2)$, and that $\delta(c_3)$ is halfway between $\delta(c_1)$ and $\delta(c_2)$. According to expected utility theory, the desirability of $(c_1, p, c_2, q, 0)$ should be halfway between the desirabilities of c_1 and c_2, equal to the desirability of c_3. However, if $\pi(\frac{1}{2}) < \frac{1}{2}$, then the desirability of the prospect according to prospect theory will be less than that of c_3. In this case the decision weight reflects a 'risk averse' attitude whereby the agent would prefer a guaranteed c_3 to a risky prospect with an expected value equal to c_3.

For our purposes, the thing to note is the close similarity between the expected utility formula for evaluating (c_1, p, c_2) and prospect theory's formula for evaluating $(c_1, p, c_2, q, 0)$. Suppose $q = \neg p$; then in both cases we're looking for a pair of functions, θ and δ, such that the value of the prospect is given by

$$\theta(p)\delta(c_1) + \theta(\neg p)\delta(c_2)$$

The difference between them is that, for expected utility theory, θ is interpreted as the agent's degrees of belief; whereas for prospect theory θ is interpreted instead as a decision weight that reflects the combination of the agent's degrees of belief and their attitudes towards risk.[4] Thus is it possible, as Kahneman & Tversky observe (1979, 280), to infer decision weights from preferences over simple prospects in a manner that's not dissimilar from how we go about inferring degrees of belief in the Ramseyan approach. Moreover, and with the appropriate additional axioms on preference, those decision weights can in turn be decomposed into a belief function and a risk function (e.g., Wakker 2004).

The end result is only a light modification on the Ramseyan theme: the degree of a belief is not quite a measure of the extent to which we are prepared to act on it, but instead a measure of the extent we're prepared to act on it *given* our attitudes towards risk. Either way, the meaning of the numerical representation of belief is manifest in the role that representation plays in a decision-theoretic context, and such representations tend to play much the same kind of role regardless of the theory. The details change, and along with them some of the axioms of the qualitative system, but in outline the general approach to explaining the quantitation of belief remains essentially the same.

[4] I'm simplifying, but only slightly. Another difference between expected utility theory and prospect theory is that the latter's weights $\theta(p)$ and $\theta(\neg p)$ needn't sum to one, so we need a more general axioms to represent prospect theory.

Bibliography

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

Alon, S. and D. Schmeidler (2014). Purely subjective maxmin expected utility. *Journal of Economic Theory 152*, 382–412.

Augustin, T., F. Coolen, G. Cooman, and M. Troffaes (Eds.) (2014). *Introduction to Imprecise Probabilities.* Maylaysia: Wiley.

Baccelli, J. (2020). Beyond the metrological viewpoint. *Studies in History and Philosophy of Science Part A 80*, 56–61.

Bolker, E. (1967). A simultaneous axiomatization of utility and subjective probability. *Philosophy of Science 34*(4), 333–340.

Borsboom, D. (2005). *Measuring the mind: conceptual issues in contemporary psychometrics.* Cambridge University Press.

Brickhill, H. and L. Horsten (2018). Triangulating Non-Archimedean Probability. *Review of Symbolic Logic 11*, 519–546.

Buchak, L. (2013). *Risk and Rationality.* Oxford University Press.

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

Bunge, M. (1973). On Confusing 'Measure' with 'Measurement' in the Methodology of Behavioral Science. In *The Methodological Unity of Science*, pp. 105–122. Dordrecht: D. Reidel Publishing.

Chalmers, D. (2011). The Nature of Epistemic Space. In A. Egan and B. Weatherson (Eds.), *Epistemic Modality*, pp. 60–107. Oxford: Oxford University Press.

Christensen, D. (2001). Preference-based arguments for probabilism. *Philosophy of Science 68*(3), 356–376.

Clark, S. (2000). The Measurement of Qualitative Probability. *Journal of Mathematical Psychology 44*, 464–479.

Davidson, D. and P. Suppes (1956). A finitistic axiomatization of subjective probability and utility. *Econometrica 24*(3), 264–275.

Davidson, D., P. Suppes, and S. Siegel (1957). *Decision making; an experimental approach.* Stanford University Press.

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

Debreu, G. (1959). Cardinal utility for even-chance mixtures of pairs of sure prospects. *The Review of Economic Studies 28*(3), 174–177.

Decoene, S., P. Onghena, and R. Janssen (1995). Representationalism under Attack: Review of An Introduction to the Logic of Psychological Measurement. *Journal of Mathematical Psychology 39*, 234–242.

DiBella, N. (2018). The Qualitative Paradox of Non-Conglomerability. *Synthese 195*, 1181–1210.

Domotor, Z. (1970). Qualitative information and entropy structures. In J. Hintikka and P. Suppes (Eds.), *Information and inference*, pp. 148–194. Reidel.

Domotor, Z. (1978). Axiomatizaton of Jeffrey Utilities. *Synthese 39*, 165–210.

Elliott, E. (2017a). Probabilism, Representation Theorems, and Whether Deliberation Crowds out Prediction. *Erkenntnis 82*(2), 379–399.

Elliott, E. (2017b). Ramsey without Ethical Neutrality: A New Representation Theorem. *Mind 126*(501), 1–51.

Elliott, E. (2017c). A Representation Theorem for Frequently Irrational Agents. *Journal of Philosophical Logic 46*(5), 467–506.

Elliott, E. (2019a). Betting Against the Zen Monk. *Synthese.*

Elliott, E. (2019b). Impossible Worlds and Partial Belief. *Synthese 196*(8), 3433–3458.

Ellis, B. (1968). *Basic Concepts of Measurement Theory.* Cambridge University Press.

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

Evren, O. and E. Ok (2011). On the multi-utility representation of preference relations. *Journal of Mathematical Economics 47*(4-5), 554–563.

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

Fishburn, P. (1967). Preference-based definitions of subjective probability. *The Annals of Mathematical Statistics 38*(6), 1605–1617.

Halpern, J. (2001). Lexicographic probability, conditional probability, and nonstandard probability. In *Proceedings of the 8th conference on Theoretical aspects of rationality and knowledge*, pp. 17–30. Morgan Kaufmann Publishers.

Hawthorne, J. (2016). A Logic of Comparative Support: Qualitative Conditional Probability Relations Representable by Popper Functions. In A. Hájek and C. Hitchcock (Eds.), *Oxford Handbook of Probabilities and Philosophy.* Oxford: Oxford University Press.

Hölder, O. (1901). Die Axiome der Quantitat und die Lehre vom Mass. *Ver. Verh. Kgl. Saschisis. Ges. Wiss. Leipzig, Math.-Physc. Classe 53*, 1–63.

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

Hájek, A. (2016). Deliberation Welcomes Prediction. *Episteme 13*(4), 507–528.

Jeffrey, R. (1965). *The Logic of Decision.* McGraw-Hill Book Company.

Jeffrey, R. (1968). Probable knowledge. *Studies in Logic and the Foundations of Mathematics 51*, 166–190.

Jeffrey, R. (1978). Axiomatizing the logic of decision. In *Foundations and Applications of Decision Theory*, pp. 227–231. Springer.

Jeffrey, R. (1990). *The Logic of Decision (Second Edition)*. University of Chicago Press.

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. (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.

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

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

Kraft, C., J. Pratt, and A. Seidenberg (1959). Intuitive Probability on Finite Sets. *The Annals of Mathematical Statistics 30*(2), 408–419.

Krantz, D., R. Luce, P. Suppes, and A. Tversky (1971). *Foundations of measurement, Vol. I: Additive and polynomial representations*. Academic Press.

Kyburg, H. (1984). *Theory and measurement*. Cambridge University Press.

Lewis, D. (1974). Radical interpretation. *Synthese 27*(3), 331–344.

Lewis, D. (1979). Attitudes De Dicto and De Se. *The Philosophical Review 88*(4), 513–543.

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

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

Luce, R. (1978). Dimensionally Invariant Numerical Laws Correspond to Meaningful Qualitative Relations. *Philosophy of Science 45*(1), 1–16.

Luce, R., D. Krantz, P. Suppes, and A. Tversky (1990). *Foundations of Measurement, Vol. III: Representation, Axiomatization, and Invariance*. New York: Dover.

Luce, R. and L. Narens (1978). Qualitative independence in probability theory. *Theory and Decision 9*, 225–239.

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

Mari, L. (2005). The problem of foundations of measurement. *Measurement 38*, 259–266.

Mari, L., P. Carbone, A. Giordani, and D. Petri (2017). A structural interpretation of measurement and some related epistemological issues. *Studies in History and Philosophy of Science 65*, 46–56.

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

Meacham, C. and J. Weisberg (2011). Representation Theorems and the Foundations of Decision Theory. *Australasian Journal of Philosophy 89*(4), 641–663.

Michell, J. (2021). Representational Measurement Theory: Is its number up? *Theory & Psychology 31*, 3–23.

Mundy, B. (1987). The Metaphysics of Quantity. *Philosophical studies 51*, 29–54.

Mundy, B. (1994). Quantity, representation and geometry. In P. Humphries (Ed.), *Patrick Suppes: Scientific philosopher*, pp. 59–102. Kluwer.

Narens, L. (1980). On Qualitative Axiomatizations for Probability Theory. *Journal of Philosophical Logic 9*, 143–151.

Narens, L. (1981). On the scales of measurement. *Journal of Mathematical Psychology 24*(3), 249–275.

Narens, L. (1985). *Abstract Measurement Theory*. MIT Press.

Narens, L. and D. Luce (1993). Further comments on the "nonrevolution" arising from axiomatic measurement theory. *Psychological Science 4*, 127–130.

Nolan, D. (1997). Impossible Worlds: A Modest Approach. *Notre Dame Journal of Formal Logic 38*, 535–72.

Nolan, D. (2013). Impossible Worlds. *Philosophy Compass 8*(4), 360–372.

Pfanzagl, J. (1968). *Theory of Measurement*. New York: Wiley.

Ramsey, F. (1931). Truth and probability. In R. Braithwaite (Ed.), *The Foundations of Mathematics and Other Logical Essays*, pp. 156–198. London: Routledge.

Reiss, J. (2016). *Error in economics: towards a more evidence-based methodology*. Routledge.

Roberts, F. (1985). *Measurement Theory with Applications to Decisionmaking, Utility, and the Social Sciences*. Cambridge University Press.

Savage, L. J. (1954). *The Foundations of Statistics*. Dover.

Scott, D. (1964). Measurement Structures and Linear Inequalities. *Journal of Mathematical Psychology 1*(2), 233–247.

Spohn, W. (1977). Where Luce and Krantz do really generalize Savage's decision model. *Erkenntnis 11*(1), 113–134.

Stalnaker, R. C. (1984). *Inquiry*. London: The MIT Press.

Stevens, S. (1946). On the Theory of Scales of Measurement. *Science 103*(2684), 677–680.

Suppes, P. (1969). *Studies in the Methodology and Foundations of Science: Selected Papers from 1951 to 1969*. Dordrecht: Springer.

Suppes, P. (2014). Using Padoa's principle to prove the non-definability, in terms of each other, of the three fundamental qualitative concepts of comparative probability, independence and comparative uncertainty, with some new axoms of qualitative independence and uncertainty included. *Journal of Mathematical Psychology 60*, 47–57. (Yes, that's actually the title).

Suppes, P. and A. Pederson (2016). Qualitative Axioms of Uncertainty as a Foundation for Probability and Decision-Making. *Minds and Machines 26*, 185–202.

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

Suppes, P. and M. Zanotti (1982). Necessary and Sufficient Qualitative Axioms for Conditional Probability. *Z. Wahrschelnllchkeltstheorle verw. Gebiete 60*, 163–169.

Suppes, P. and J. Zinnes (1963). Basic Measurement Theory. In D. R. Luce (Ed.), *Handbook of Mathematical Psychology*. John Wiley & Sons.

Swoyer, C. (1991). Structural Representation and Surrogative Reasoning. *Synthese 87*(3), 449–508.

Titelbaum, M. (2022). *Fundamentals of Bayesian Epistemology 2: Arguments, Challenges, Alternatives*. Oxford University Press.

Wakker, P. (2004). On the composition of risk preference and belief. *Psychological Review 111*, 236–241.

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

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