

2. Subjective Probability and the Content/Attitude Distinction

Jennifer Rose Carr

On an attractive, naturalistically respectable theory of intentionality, mental contents are a form of measurement system for representing behavioral and psychological dispositions. This chapter argues that a consequence of this view is that there is substantial arbitrariness in the content/attitude distinction. Whether some measurement of mental states counts as characterizing the *content* of mental states or the *attitude* is not a question of empirical discovery but of theoretical utility. (By *theoretical utility*, I mean practical utility associated with theory adoption that does not derive from the value of accuracy.) If correct, this observation has important ramifications in a number of philosophical literatures, not just in philosophy of mind but also in epistemology and rational choice theory. The focus of this chapter is on the relation between the measure theory of mental content and the modeling of subjective probability in doxastic states.

Subjective probabilities are expressed by statements like “It probably rained in Leeds last night” or “Ten to one he’s going to lose his temper.” Beyond that, there is considerable controversy. There are a variety of competing views about how to understand and model the doxastic states associated with subjective probability. The three most popular views model subjective probabilities as (i) binary beliefs about probability, (ii) precise credences, and (iii) imprecise credences. On the binary belief view, an agent’s total doxastic state is best modeled as a function from propositions to one of three values, corresponding to belief, suspension of judgment, and rejection. On the precise credence view, an agent’s total doxastic state is best modeled as a credence function: a function from propositions to real numbers, where greater numbers represent greater certainty in the relevant proposition. On a common version of the imprecise credence view,¹ an agent’s total doxastic state is best modeled as a set of credence functions.²

¹ There are a variety of alternative imprecise credence views. For example, Kyburg (1983) defends an imprecise model whereby an agent’s doxastic states are represented by a function that maps propositions to interval subsets of $[0,1]$; Sturgeon (2008) defends a model whereby imprecise credences can also exhibit a form of vagueness at their borders that is not captured in either the set of probabilities model or the interval model.

² These options are not exhaustive. One alternative is the view that an agent’s total doxastic state is best represented by a preorder over propositions, representing comparative confidence without cardinal degrees of confidence (Keynes 1921; de Finetti 1937; Koopman 1940; Savage 1972; Fine 1973). Another, more common alternative is the view that agents have both binary

36 | Jennifer Rose Carr

These different modeling claims are generally thought to correspond to claims about the fundamental descriptive nature of doxastic states. Are doxastic attitudes toward a proposition p coarse-grained, such that an agent either believes p , rejects p , or suspends judgment about p ? Or are these attitudes more fine-grained? Are they fine-grained and sharp, like the maximum horizontal width of Chile, or fine-grained and imprecise, like the latitude of Chile?

These are non-normative questions. But the different belief models are often linked with theories of rationality: for example, that it's irrational to have imprecise credences³ or irrational to have precise credences.⁴

The theory of intentionality and the theory of rationality are typically treated as orthogonal.⁵ But this chapter argues that an attractive theory of intentionality has important ramifications for rationality. If the measure theory of mental content is correct, I argue, then it cannot be the case that agents are irrational by virtue of having probabilistic binary beliefs, or having precise credences, or having imprecise credences, in response to their overall epistemic situation.

The central argument of the chapter goes as follows:

- (1) If the measure theory of mental content is true, then the content/attitude distinction is measurement system relative.
- (2) If the content/attitude distinction is measurement system relative, then there's no psychological difference between having probabilistic binary beliefs, precise credences, and imprecise credences.
- (3) If there's no psychological difference between having probabilistic binary beliefs, having precise credences, and having imprecise credences, then none of these is rationally impermissible.
- (C) So, if the measure theory of mental content is true, then it is not rationally impermissible to have probabilistic binary beliefs, or precise credences, or imprecise credences.

First clarification: I use "psychological" to isolate those features of an agent's mental states that are either intrinsic to the agent or grounded in the

beliefs and some form of degrees of belief, precise or imprecise. How these two relate to each other is a topic of debate. Some hold that our doxastic states should be modeled with both precise and imprecise credences: that though agents have imprecise credences, there may be some precise credence function at a time that an agent most identifies with and that explains her behavior (Moss 2015). Some even hold three different models are independently needed: belief, comparative confidence, and some form of credence model (Fitelson manuscript). The conclusion of this chapter generalizes to these other belief models, though the details of the generalization require care.

³ See e.g. White (2009); Elga (2010).

⁴ See e.g. Walley (1991); Joyce (2010).

⁵ One obvious exception involves theories of intentionality that make use of possible worlds propositions, which treat all necessarily true propositions as equivalent. (Likewise for all necessarily false propositions.) Plausibly, it is irrational to believe (or have high credence in) p while suspending judgment about (or having non-low credence in) $\neg p$. But arguably, it's not irrational to be uncertain about some abstruse mathematical falsehood while believing that it's either raining or not raining.

interactions between the agent and her environment. (This is intended to be neutral with respect to internalism versus externalism about mental content.) What I mean to rule out are those features of the agent that depend on how the agent is described by the theorist: for example, whether she's described in English or Zazaki, or whether she's described in inches or centimeters, or whether she's described in the precise credence framework or the imprecise credence framework.

Second clarification: the consequent of (3) and of (C) is shorthand for the claim that, for any epistemic situation, an agent cannot be irrational merely by virtue of having probabilistic binary beliefs, nor merely by virtue of having precise credences, nor merely by virtue of having imprecise credences. (An *epistemic situation* specifies the conditions that are relevant for determining what total doxastic states are rationally permissible. According to evidentialists, this amounts to the agent's evidence. For others, other factors are relevant.) I mean to claim something stronger than that for each of the three doxastic types under discussion, there's some epistemic situation such that that doxastic type is permissible.

The chapter will proceed as follows: in section 1, I'll spell out the measure theory of mental content, first motivating the theory in 1.1 and then giving a toy measure theory for belief and desire in 1.2. Section 2 defends the first premise of the argument above. Section 3 defends the second premise, first comparing probabilistic binary beliefs with credences in 3.1 and then comparing precise with imprecise credences in 3.2. I close in section 4 with a brief defense of the third premise.

1. Content

1.1. *The Measure Theory of Mental Content*

Here's one picture of how minds have content: there is some non-physical relation between physical states of brains and certain abstract objects (propositions). By virtue of standing in this non-physical relation, creatures with brains are able to have contentful mental states. These relations are *discovered* when theorists observe agents' behavior, transitions among physically realized functional states of brains, and whatever else might form the physical basis for mental states.

Here's an alternative: the relation between mental states and abstract contents is not discovered but *constructed*. "Attitudes"—in particular, doxastic and orrectic states—are just clusters of dispositions. They are not clusters of dispositions that also happen to have a content by virtue of standing in a special non-physical relation to a proposition.

On this latter theory, attitudes are related to propositions in the way that temperatures, lengths, and weights are related to numbers. What does it mean to say that a particular pizza is a fourteen-inch pizza, or that a particular beer stein holds thirty-four fluid ounces? Does the pizza have that diameter in

virtue of standing in a non-physical relation to the number fourteen? If this relation exists, it is neither mysterious nor explanatory. We know that the relevant relationship between the pizza and the number fourteen is constructed. Physical objects possess intrinsic and nonintrinsic physical properties—diameter, capacity, etc.—that are compactly represented by indexing them to representatives. Abstract objects such as numbers are well-suited for this purpose, for reasons we'll see.

So similarly, physical objects with minds possess intrinsic and nonintrinsic physical properties—dispositions to form mental states, to act, and so on—that are represented by indexing them to abstracta like propositions.⁶ This view, which has been proposed and rediscovered many times over the past six decades,⁷ is sometimes called the “measure theory of mind.”⁸

Imagine a time before numerical temperature scales have been created. Theorists understand that objects can be colder, hotter, etc. There is an infinite range of mutually exclusive possible thermal properties an object can have. It would be highly inconvenient to assign each such property a different word. We could describe them compositionally: *really hot, kind of hot, warm, . . .* But this wouldn't give us as much detail as we might find useful.

We notice that there are certain structural relations between the various possible thermal properties of objects. Objects can be totally ordered by hotness; some object could be a little hotter than another, or a lot; for any two thermal properties, there is a thermal property midway between the two; and so on. We notice that the space of numbers also has these structural relations. So to characterize the range of mutually exclusive possible thermal properties, it's convenient to index different thermal properties to real numbers. Thus a tepid object must be assigned a number that's somewhere between the numbers assigned to a hot object and a cold object; a warm object's number must be between that of a hot object and a tepid object; and so on. Which precise numbers we use for which temperatures is arbitrary.

What is non-arbitrary is the relations amongst these numbers: they are assigned so as to preserve certain structural relations among the relevant thermal properties. For example, relative to a shared scale, the number representing the temperature of x should be greater than the number representing the temperature of y iff x is hotter than y . Of course, not all relations (or even all natural or interesting relations) between real numbers need to represent corresponding relations among thermal properties in order for numbers to be usable for representing thermal properties. We needn't say that a day when the temperature's 80°F is twice as hot as a day when it's 40°F, even though

⁶ This picture is one of many, many theories of mental content. Others include Stampe (1977); Dretske (1981); Fodor (1987); Millikan (1989a, 1989b); Kriegel (2013). This survey is highly partial (in both senses). I will not address the many merits and shortcomings of alternative theories.

⁷ Versions of this suggestion have been discussed and defended by Suppes and Zinnes (1963); Churchland (1979); Field (1980); Stalnaker (1984); Dennett (1987); Davidson (1989); Matthews (1994).

⁸ Matthews (1994).

80 and 40 are related in that way. Indeed, what relations are relevant for the representation of thermal properties isn't fully determined by facts about thermal properties.⁹

Next: imagine a time before propositional attitude ascriptions. Agents have complex dispositions to produce various forms of behavior, transitions between brain states and functional states, and so on. We notice that these dispositions have certain structural relations. Some are incompatible (rationally or psychologically); some necessitate others (rationally or psychologically); some are orthogonal; etc.

We can represent these structural relations by indexing agents' psychological dispositions to abstract objects that have analogous structural relations. Here, real numbers are not the obvious choice. Numbers are a natural choice for many measurable attributes because of their ordinal and cardinal properties: temperatures, weights, lengths, and so on are conveniently represented with reals in part because an object can be a little hotter/heavier/longer/etc. than another, or a lot hotter/heavier/longer/etc. But typically we don't have much use for ordering psychological dispositions. The relations among psychological dispositions that we happen to be interested in representing are relations of incompatibility, necessitation, orthogonality, etc. So *sets* are convenient for representing psychological dispositions. If two dispositions are rationally incompatible, e.g., we can capture this by indexing them to disjoint sets.¹⁰

So there is some analogy between the relation between thermal states and numbers and the relation between mental states and propositions. How deep does this analogy go?

The reason why thermal properties are conveniently represented with real numbers is that the range of possible mutually exclusive thermal properties shares certain structural properties with the set of real numbers. But again, how these real numbers are mapped to different thermal properties is largely arbitrary: there are different temperature scales that preserve the relevant ordinal and cardinal relations between numbers and thermal properties. The scales that we in fact use are convenient for various purposes—using 0 and 100 to privilege the phase change temperatures for water at sea level, or to privilege really cold and really hot air temperatures in Western European climates, etc.—but the selection of scale is fundamentally conventional. There are infinitely many different temperature scales that we could have used with equal fidelity to the physical facts.

There are more and less conservative interpretations of the measurement analogy.

⁹ I will not rehearse the details of standard expositions of measurement theory. For classic expositions, see Suppes and Zinnes (1963) and Krantz and Tversky (1971).

¹⁰ Psychological states' relations of incompatibility, entailment, etc. could also be representable with various forms of structured propositions or even (with some work) sentences in artificial or natural language. (See Matthews (1994).)

40 | Jennifer Rose Carr

- *Conservative*: the relation between mental states and abstract objects is illuminated by, and no more mysterious than, the relation between thermal properties and abstract objects.¹¹
- *Less conservative*: content ascriptions are literally measurement reports. As such, certain features of the representation space of indices are arbitrary. There are multiple equally adequate measurement “scales.”
- *Even less conservative*: not only are there different possible “scales” for measuring psychological states, but there are also different psychological properties that might be useful to measure. None has decisive claim to the title “content.”

On the least conservative version of the view, propositional contents are nothing more than tags that theorists use to keep track of subjects’ behavioral and psychological properties. For a state to have propositional content is for it to be indexed relative to some representation space of propositions. Different kinds of objects (structured or unstructured, abstract or non-abstract) can form a representation space for mental states. This is the version of the analogy that we’ll explore, and to which we’ll apply the title “measure theory of mental content.”

1.2. A Toy Measure Theory

Mental states have a variety of different properties that could be useful to index within some representation space for the purpose of measurement. Here is an oversimplified example of how properties of mental states motivate the use of a particular kind of representation space of abstracta:

First, we specify which properties of an agent’s mental states are relevant to our interests. Suppose these are twofold: the mental states’ rationality conditions and their causal tendencies.¹² To capture the interesting relations amongst attitudes, and between attitudes and their environment, we can index them to abstract objects that are interrelated in parallel ways.

We assume that agents with mental states are, as a noncontingent matter, largely rational.¹³ So the rational interrelations between mental states are psychologically noncontingent. We then observe that *set theoretic* or boolean relations can conveniently represent rational relations between mental states. How? We can index each theoretically interesting mental state to some set. If one mental state is rationally or psychologically incompatible with another, we

¹¹ This seems to be the version of the thesis endorsed in Matthews (1994).

¹² There are various possible motivations for this choice of properties to index to “content.” The assumption that agents must be (widely) rational in order to have contentful mental states is widespread but not universal. It is sometimes tied to the radical interpretation view (Davidson 1973; Lewis 1974; Williamson 2009, ch. 8). The relevance of causal information is also widespread (from Fodor (1987) to Millikan (1989a, 1989b) and Dretske (1988).) The view I sketch here is a simplified version of the view defended in Stalnaker (1984, ch. 1).

¹³ We’ll take no stand on which kinds of entities will count as agents; but if a cat or an insect counts as an agent, then it is largely rational.

make sure to index them to disjoint sets. If one mental state rationally or psychologically necessitates another, we make the first a subset of the second. If the two mental states are rationally or psychologically compatible but neither necessitates the other, we index them to overlapping sets such that neither is a subset of the other. And so on.

In principle, this could be done with any sets of large enough cardinality: sets of numbers, pizzas, or stars. But we are also interested in the causal relations between attitudes and their environment. So it's more theoretically valuable to use sets that clue us in to these causal relations. For example, suppose we've already divided mental states on the basis of functional role such that one subset contains beliefs and another contains desires.¹⁴ We index each belief to a set of possible worlds—the worlds where the facts¹⁵ that tend to cause the belief obtain.¹⁶ We index each desire to a set of possible worlds: the worlds where the facts that the desire tends to cause obtain.

Now, from knowing that some belief is indexed to a certain set of worlds, we are able to read off certain facts about the state: that the agent can't rationally hold a belief indexed to any set disjoint from that set, for example, and that the belief is one that tends to be caused by facts that obtain in worlds in that set. So, in our toy model, if a belief tends to be caused by the fact that p is true, then we index the belief with the possible worlds proposition p ; we say that the belief has p as its content. If a desire tends to cause q to be true, then we index the desire with the possible worlds proposition q ; we say that the desire has q as its content.

The representation space of sets of possible worlds can play the two roles we'd aimed for: set relations mirror rational relations; sets of possible worlds encode the tendential causal relations between mental states and their environments.

This is obviously a simplification. We might, for example, be interested in representing more features of mental states than their rational interrelations and causal relations with their environment. If so, there are decisions to be made about whether, in representing these other features, it is more useful to use a separate representation space (treating these as distinct measurements), or a unified representation space. What motivates choosing some specific representation space as the space of *contents*?

2. Relativism about the Content/Attitude Distinction

Clearly, even in our toy measure theory, some properties of mental states are not reconstructable from the propositions to which they are indexed. Not everything is built into the content. A belief and a desire may have the same propositional content, but they are not the same mental state. Some

¹⁴ I remain neutral about how this might work; the story we tell is consistent with indication theories, teleosemantics, and so on.

¹⁵ For ease of exposition, I assume that the relevant causal relata are mental states and facts.

¹⁶ Simplifying wildly, I elide complexities introduced by false belief. These could be filled in using Dretske's "channel conditions," Stalnaker's "normal conditions," and so on. I also ignore distal versus proximal causes.

properties of mental states are encoded in propositional contents, while other properties of mental states are not. These other properties are relegated to the “attitude,” “form,” “vehicle,” etc.

We can encode properties of mental states in various ways. There are different properties of an agents’ total set of psychological and behavioral dispositions that can be measured using different systems of indices. For some such properties, we can conveniently use real numbers. Obvious candidates for this treatment are credences and utilities. For others, features of mental states, we can use sets, or sequences of objects, properties, and relations, or sentences, or . . . These different indexing tools can also be combined.

Return to our example motivations for a representation space of sets of possible worlds. The causal tendencies linking mental states and their environment are plausibly vague, and come in degrees: hence credences and utilities. But if we think of attitudes as coming in degrees, can we retain the same rationale for using sets of worlds as contents? For binary beliefs, being indexed to disjoint sets indicated rational or psychological incompatibility. But one can rationally have exactly the same degree of belief in both $\neg p$ and $\neg\neg p$: namely, credence 0.5. So disjoint sets as contents no longer represent rational or psychological incompatibility.

Similarly, what if we want to represent features of attitudes other than causal relations? We often take some of our beliefs to be “about” things that arguably couldn’t stand in causal relations to them, e.g. mathematical facts, normative facts, and so on. So content ascriptions aren’t confined to encoding information about the causal conditions and causal tendencies of mental states.

In both cases, the measure theorist faces the question: should we represent the relevant features of mental states within the representation space that we use for attributing mental *content*? Or should we separately specify these features in our characterization of the mental state as part of the *attitude*?

The measure theorist is committed to the conclusion that answering these questions might be a matter of decision, not discovery.

There are (at least) three ways in which the choice of measurement system is susceptible to arbitrariness¹⁷:

1. There is arbitrariness in the assignment of indices to physical properties.

In quantitative measurements, this affects choice of scale. We measure thermal properties with numbers, but the assignment of numbers is conventional and arbitrary (up to the point of preserving the relevant structural relations of thermal properties within the representation space). There are different temperature scales, none of which can be sensibly thought of as the “one true temperature scale.” Any particular temperature scale is at most preferable

¹⁷ That is, arbitrariness with respect to the goal of accurately representing psychological states. Pragmatic considerations may motivate a choice of measurement system; alethic considerations cannot.

within particular contexts, relatively to local conventions or other sources of practical utility.

2. There is arbitrariness in the choice of which physical properties to measure.

For ordinary objects, it's obvious that there is a huge variety of physical quantities that are worth measuring in different contexts, for different purposes. You might be most interested in the pizza's diameter, because you're hungry, while I might be most interested in its temperature, because I want to avoid cheese-burn, while the delivery driver might be most interested in its speed, to ensure it arrives in thirty minutes or less. None of these systems of measurements is privileged (capturing the "numerical content" of the pizza); which system we should use depends on our purposes, what features of objects we want to keep track of.

3. There is arbitrariness in the individuation of measurements.¹⁸

We can measure an object's speed; we can measure its direction; we can measure its temperature. We can also measure its velocity = ⟨speed, direction⟩; we could, if we wanted, measure its ⟨temperature, direction⟩. Obviously, it's not the case that *speed* is a defective system of measurement because it doesn't encode information about direction. It just measures a somewhat different physical quantity from what velocity measures. Should we use a velocity measurement, or two separate measurements of speed and direction? Both do equally well with respect to information about the physical facts. Similarly, we might describe an object's movement in terms of its constant acceleration over an interval of time, or we might describe the same movement in terms of its speed at each of the times within the interval. Whether we count this as one measurement or many makes sense only relative to a choice of measurement system. The choice of measurement system we use is a matter of theoretical utility rather than accurate representation of the physical facts.

If the measure theory of mental content is correct, then we should expect each of these to hold for mental content attributions.

First, we should expect there to be multiple adequate representation spaces of contents, which are more or less preferable relative to the theorist's aims in a context, and not relative to "the one true assignment of contents."

Second, creatures like us have a huge variety of potentially interesting psychological properties. We should expect that there may be no unique privileged body of content-determining properties.

Third, we should expect that, insofar as different measurement systems individuate the psychological properties relevant for contents in different ways, there will be arbitrariness in content individuation: relative to different measurement systems, the same mental state might count as one attitude toward one proposition or multiple attitudes toward multiple propositions.

¹⁸ Here, "measurements" does not mean acts of collecting information about the system ("taking measurements,") but rather the unit-sized chunks of information so acquired.

One might respond that the first observation can be settled by stipulation: the *propositional* indices will count as content. But which indices are propositional? Because there are several live hypotheses about the nature of propositions, this is unhelpful.¹⁹ Moreover, the stipulation is not innocuous: expressivists often view it as a central thesis of their positions that there is a type of judgment, expressible in assertoric sentences, that has non-propositional content.²⁰ Similarly, many have argued that perceptual experiences have non-propositional content.²¹ Finally, as we'll see, whether some property of a mental state should feature in the proposition to which it indexed or some other measurement is not settled by facts about the agent's dispositions, brain states, causal history, environment, etc. If the measure theory is correct, then this question is only answerable relative to a choice of measurement system. In a slogan, the measure theory is committed to mental content relativism:

Mental content relativism: What content(s), if any, a psychological state has is measurement system relative.

More importantly for present purposes, the measure theory is committed to mental content/attitude relativism:

Mental content/attitude relativism: Whether some property of a psychological state is part of the content or the attitude is measurement system relative.

These forms of relativism should be distinguished from *antirealism about content*. I won't attempt the foolhardy task of giving a general definition of antirealism. I'll merely note that you don't count as an antirealist about whether a bird is flying at twenty mph merely because you concede that you could have given its speed in km/h, or given its velocity instead. Suppose the sentence "The bird is flying at twenty mph" is felicitously uttered. The proposition expressed by that utterance will be true or false depending solely on the motion of the bird.

Similarly, you don't count as an antirealist about whether a mental state is a belief that *p* (relative to a representation space) merely because you concede that it could have been equally well characterized in a different measurement system, or you could have described some mental state constituted by an overlapping but slightly different cluster of dispositions, physical states, etc.²²

Similarly for *instrumentalism about content*, the view that propositional attitude ascriptions are not literally true (or truth-evaluable), and are merely useful devices for specific purposes. Measurements are literally true. It's just that which things we choose to measure, and which units we use, depend on our purposes.

¹⁹ Propositions may be complexes of senses (Frege 1892), or complexes of objects, properties, and relations (Russell 1903), or sets of possible worlds (Stalnaker 1984), or cognitive event types (Soames 2014), or . . .

²⁰ See, e.g., Gibbard (1990, 2003); Schroeder (2009).

²¹ e.g. Tye (2000).

²² The view is, however, antirealist about the content-attitude distinction, in the sense that this is merely a feature of measurement systems for characterizing psychological states.

Objection. The analogy between content and speed breaks down in contexts of disagreement. (4a) does not conflict with (4b):

- (4) a. The bird is flying at 20 mph.
- b. The bird is flying at 32 km/h.

But (5a) does conflict with (5b):

- (5) a. Property F features in the content of state s .
- b. F does not feature in the content of s .

A realist about these claims is committed to the hypothesis that one is false. The relativist²³ sidesteps this by treating (5a) and (5b) as semantically incomplete. They are only interpretable as true or false if relativized to a measurement system:

- (6) a. Property F features in the content of state s according to measurement system M_1 .
- b. F does not feature in the content of s according to measurement system M_2 .

Ordinary mental content attributions do not explicitly relativize to measurement systems and so have the form of (5a) and (5b), which according to relativism are semantically incomplete. So relativism about content and about the content/attitude distinction is in fact antirealist about ordinary mental content attributions.

Reply. On the relativist view, the claim that s has F must be relativized to a measurement system in order to express a truth-evaluable proposition. But there's no clear reason for requiring this relativization to be explicit. It doesn't seem to me that ordinary speakers exhibit semantic blindness or make some other kind of mistake of linguistic competence by virtue of the fact that they don't explicitly relativize their mental state attributions to a measurement system. The selection of measurement system may sometimes be provided by some relevant context.²⁴ Compare: given special relativity, claims about (e.g.) length must be evaluated relative to a reference frame. In accepting special relativity, we are not thereby committed to antirealism about length. Similarly, sentences that include non-anaphoric pronouns and demonstratives must be evaluated relative to an assignment of referents to the free variables. But we needn't be antirealists about facts expressible in indexical language.²⁵

²³ In the vein of Harman (1996).

²⁴ This view is also meant to be neutral with respect to relativism about truth in the vein of MacFarlane (2014), according to which sentence truth is relativized to both the context of utterance and a context of assessment; cf. Egan (2007).

²⁵ Caveat: the terms "antirealism" and "relativism" are precisified in a variety of ways in a variety of philosophical literatures. There are some precisifications that are incompatible with my usage. My aim is to situate the view under discussion by contrasting it with alternative possible views, not to defend it against the charge of antirealism under any interpretation of the word. Thanks to Jonathan Livengood for pressing this objection.

46 | Jennifer Rose Carr

Relativism about the content/attitude distinction has an affinity with interpretivism:

Interpretivism: an agent's (partial or binary) beliefs and desires are whatever beliefs and desires an ideal interpreter would ascribe to the agent in order to make the best (explanatory and rationalizing) sense of the agent's dispositions to act.²⁶

Interpretivism of some form or another has been a highly influential view among formal epistemologists and decision theorists. Understanding contents as measurements provides a kind of answer to a challenge for interpretivists: why are beliefs and desires the kinds of things that necessarily coincide with the ideal interpreter's best interpretation? On the view under discussion, to have a contentful attitude is to have a cluster of properties that an (in some sense) ideal interpreter might index to some kind of object that can enter into boolean relations. Contents just are conventional tags that are useful in interpretations.²⁷

Note, however, that the two views are, strictly speaking, orthogonal. An interpretivist can accept that there's a unique, privileged space of contents that agents have attitudes toward, and hold that ideal interpretation assigns attitudes to these contents.²⁸ And one can hold that there's arbitrariness in the content/attitude distinction while denying that mental states have an illuminating relationship with ideal interpreters. (What speed a car is going at has no illuminating relationship with the speed attributed to it by an ideal interpreter.)

3. Ramifications for Belief Modeling

3.1. *Binary Belief vs Credence*

Some have denied that ordinary agents have credences, precise or imprecise.²⁹ There are a variety of reasons for this conclusion, often having to do with computational limitations and the theoretical superfluity of quantitative attitudes. We may have some attitudes toward probabilistic contents, these philosophers argue, but not all of our doxastic states have a tacit commitment to probabilities.

A general challenge for reducing credences to binary belief is to account for subjective probability. Suppose we take it as a datum that I assign subjective probability 0.5 to the proposition that a coin will land heads the next time it's

²⁶ See in particular Lewis (1974).

²⁷ Relativism about the content/attitude distinction suggests that what counts as a best interpretation depends partly on the interpreter's purposes. A result of this view is that there may not be a unique privileged interpretation, for reasons that are totally independent of indeterminacy of reference.

²⁸ Williams (manuscript) argues that those who defend probabilism in terms of representation theorems (e.g. Savage 1954) may be committed to this view.

²⁹ Harman (1986), Holton (2015).

tossed. How can this be cashed out as a binary belief? It's clear that we have doxastic attitudes toward claims about probabilities: I might believe that the objective chance that the coin will land heads is 0.5. But binary beliefs about objective chances are not the same thing as subjective probabilities. I just as well might suspend judgment on the proposition that the objective chance that the coin will land heads is 0.5. I might be more or less confident of this proposition. Indeed, suppose I know that either the coin is biased 3-to-1 toward heads or it's biased 3-to-1 toward tails, but I'm uncertain which. I fail to believe that the objective probability of heads is 0.5—I think that the objective probability is either 0.75 or 0.25—but nevertheless my subjective probability that the coin will land heads is neither. Indeed, it might be 0.5. One can run a similar argument against the hypotheses that subjective probabilities are beliefs about evidential probability. These forms of belief about other kinds of probability are orthogonal to subjective probabilities.

Holton (2015) and (p.c.) denies that ordinary agents have credences, but accepts that we do have subjective probabilities. We don't need to appeal to other kinds of probability to explain the phenomenon of subjective probability, on Holton's view. Instead, if there are subjective probabilities, we should think of them as figuring into the content of doxastic states, rather than the attitude.

Initially, this suggestion may seem puzzling, since the notion of subjective probability is typically identified with the notion of a credence. The two phrases are sometimes taken to be synonymous. It would make no sense to reduce credences to beliefs about credences. But on Holton's view, the two notions come apart. Subjective probabilities—estimates of a primitive likelihood—are the kinds of things that people can have beliefs about.³⁰

So (going beyond Holton's view), instead of sets of worlds as contents, the contents of binary beliefs might be sets of world–probability function pairs, where the second parameter represents a possible subjective probability function. An agent's total doxastic state could be represented as a set of $\langle w, Pr \rangle$ pairs. Where the agent is committed to the probability of some possible worlds proposition p being n , all of these pairs contain a Pr coordinate such that $Pr(p) = n$. However, the agent need not assign subjective probabilities to all propositions. Where the agent withholds judgment about the probability of some possible worlds proposition p , the agent's doxastically possible $\langle w, Pr \rangle$ pairs will contain Pr s that assign different values to p .

But what is the psychological difference between a binary belief that assigns p subjective probability n and a credence that assigns p subjective probability n ? The measure theorist should say: nothing. They might involve precisely the same dispositions to cause and be caused by other mental states; they might be grounded in the same evidence and they might cause precisely the same dispositions to act. Decision rules associated with credences have notational variants expressible in terms of binary belief. If the agent's belief

³⁰ See also Sepielli (2012) on the notion of minimal probability.

state fails to determine a unique probability function, no problem: the agent's psychological dispositions will be easily representable as imprecise credences. If an agent's belief set B is a set of $\langle w, Pr \rangle$ pairs, her imprecise credal state $C = \{Pr : \exists w. \langle w, Pr \rangle \in B\}$. Here, again, whatever the correct decision norms for imprecise credences are, they can straightforwardly be generalized to binary beliefs about uncertain subjective probabilities.³¹

Here, the measure theorist should accept: whether the property of a doxastic state that we characterize as subjective probability is indexed as part of the content, or is indexed separately as a different feature of the attitude, is largely arbitrary. If it is common ground that doxastic states can have associated subjective probabilities, we can encode this information in various ways. Fans of credences measure possible worlds content and confidence separately, while fans of binary belief measure them together. So similarly, we might separately represent measurements of an object's speed and of its direction, or we might represent these as a single measurement of velocity. It is not worth debating which of these measurements captures the objective "movement content" of moving objects.

Subjective probabilities may or may not be roughly introspectable. But whether they are a feature of the content or the attitude is not. If measure theory is correct, then there's nothing there to introspect.

3.2. *Precise vs Imprecise Credence*

Relativism about the content/attitude distinction also entails that, as a psychological matter, there's no non-relative answer to the question of whether an agent's credences are precise or imprecise.

Imprecise credence attributions are typically motivated in three ways: psychologically, pragmatically, and epistemically. Briefly, these views claim:

1. *Imprecise credences are irrational but psychologically realistic*: the imprecise credence model provides a representation of uncertainty that's more descriptively plausible than attributions of precise credences.
2. *Imprecise credences are rationally permissible*: the imprecise credence model provides a representation of uncertainty that seems to correspond to certain forms of practically rationally permissible dispositions to act, which the precise model (paired with traditional decision rules) falsely claims are practically irrational.
3. *Imprecise credences are rationally required*: the imprecise credence model provides a representation of an epistemically rationally mandatory

³¹ So, for example, Bayesian decision theory requires selecting options that maximize expected utility, where the relevant expectation is calculated by the agent's credence function. In binary belief terms, we can calculate expectations relative to the probability function(s) present in the agent's belief set. If the agent's belief set fails to determine a unique probability function, the relevant set of probability functions can be used with variations of decision rules for imprecise credences: for example, E-permissibility (Joyce 2010) or α -maximin (Seidenfeld 2004).

Subjective Probability and Content/Attitude | 49

response to evidence that is ambiguous (i.e. that points in different directions) or unspecific (i.e. that doesn't point in any direction).

On the first view, precise credences are unrealistic as a representation of our doxastic states. This might be motivated by the observation that we rarely, and plausibly never, have introspective access to precise credences.³² It might be motivated by the fact that we have finite brains, incapable of infinitely sharp credences.³³

It's worth noting, however, that neither of these worries is adequately answered by the traditional imprecise credence model. The traditional model represents an imprecise credal state with a set of credence functions. This set of credence functions determines, for each proposition that the agent has a doxastic attitude toward, a set of real numbers representing the agent's imprecise credence in that proposition. This set of real numbers has infinitely sharp boundaries, which are no more introspectable than precise credences. Similarly for the conception of imprecise credences as merely assigning intervals or upper and lower previsions.³⁴

Further motivations come from the hypothesis that imprecise credences are manifested in behaviors that are treated as irrational within traditional expected utility theory. The claim that we in fact display these behaviors and thereby manifest imprecise credences is orthogonal to the claim that these behaviors are irrational, but because our present purposes are descriptive, the two views can be discussed together. Some behaviors are often associated with imprecise credences: for example, having distinct buying and selling prices for gambles,³⁵ or willingness to forgo certain forms of sure gain in diachronic betting contexts.³⁶ But these associations are based on specific assumptions about how precise credences must be manifested in behavior: specifically, that agents with precise credences are expected utility maximizers. (Neither of these forms of behavior is consistent with expected utility maximization, given fixed utilities.)

The link between a particular means of indexing psychological dispositions and rules of rational decision-making is weak. Alternative decision rules might be implemented that allow for an agent's choice behavior to be explained

³² Cases where agents have access to precise objective chance information seems to be the best candidate, but realistically these cases never arise. We never in fact learn with certainty that the coin is perfectly fair. Credence 1 and 0 might be the exception. But even here, there are worries: ordinary agents typically refuse to bet their lives on any proposition, no matter how obvious.

³³ Note: it doesn't follow from the finitude of our brains that we can't have infinitely precise credences. Finite creatures can, after all, have infinitely precise blood-alcohol levels.

³⁴ See Kyburg (1983); Pedersen and Wheeler (2014). This concern is addressed in Sturgeon (2008), who suggests that credences be both imprecise and vague. How precisely this vagueness is to be modeled is an open question. I don't consider this model further in this chapter, but see Carr (manuscript) for discussion.

³⁵ e.g. Walley (1991).

³⁶ See Elga (2010).

equally well in terms of precise credence. For example, suppose an agent has distinct buying and selling prices for gambles. These buying and selling prices for gambles might be separable and determined by (precise) higher-order probabilities. For example, an agent might have a precise credence n_i in p , but have positive credence in the proposition that the probability of p is instead n_j , and in the proposition that the probability of p is instead n_l , etc.³⁷ Let the minimal probability for p that the agent assigns positive credence to be n_1 and the maximum be n_k . We might further allow that a rational agent might bet, not in accordance with her first-order credence in p , but instead ~~with~~ in accordance with the probabilities for p she thinks are possible. She might buy bets on p as though she had precise credence n_1 and sell bets on p as though she had precise credence n_k .

Similarly, depending on the choice of decision rule, imprecise credences might be manifested in behavior that is indistinguishable from expected utility maximizing (or other behaviors associated with precise credences). For example, an agent's buying and selling prices might be determined by a privileged credence function in her representor with which the agent identifies,³⁸ or the midpoints of upper and lower credences,³⁹ etc.

If the measure theory is correct, then, there need be no psychological difference between precise and imprecise credences. These may simply be alternative measurement systems for characterizing the same set of complex psychological dispositions. Again, one represents a certain openness about probabilities in the attitude (an imprecise credence), while the other represents this in the content (uncertainty in probabilistic propositions).

These different systems have their own upsides and downsides. For example, the imprecise credence representation is less simple than the precise credence when it comes to the characterization of mental states. But if having distinct buying and selling prices for gambles is rationally permissible, then the imprecise credence representation compensates for this shortcoming in its comparatively simpler decision rule. Content/attitude relativism has no bearing on whether behaviors associated with imprecise credences are rationally permissible. But it does have bearing on whether imprecise credences are rationally required.

So we should consider the third motivation for imprecise credences, which claims that they are epistemically required as a response to ambiguous or

³⁷ What probability is at issue here? Objective and epistemic probabilities are obvious candidates, as well as hypotheses about the agent's own first-order credences. Some other conception of subjective probability, along the lines discussed in the previous subsection, is even an open possibility.

³⁸ See Moss (2015).

³⁹ As Moss notes, on the midpoint view, an agent may behave as though her credences are not probabilistic, since the midpoints of imprecise credences determined by a set of probability functions need not itself form a probability function. But this is consistent with the agent's having precise, non-probabilistic credences.

unspecific evidence. The standard example motivating rationally required imprecise credences involves contrasting two evidential situations:

Fair Coin

You have a coin that you know to be perfectly fair. What should your credence be that when you next toss the coin, it'll land heads?

Mystery Coin

You have a coin that was made at a factory where they can make coins of pretty much any bias. You have no idea whatsoever what bias your coin has. What should your credence be that when you next toss the coin, it'll land heads? (See e.g. Joyce 2010.)

There is a sharp credence that stands out as a natural candidate for both cases: 0.5. In the Fair Coin case, credence 0.5 is required by the Principal Principle. In the Mystery Coin case, you have no more reason to favor heads than to favor tails; your evidence is symmetric. Credence 0.5 in both heads and tails seems *prima facie* to preserve this symmetry.

But the proponent of imprecise credences claims this reasoning doesn't properly distinguish the evidential situation of someone in the Mystery Coin case from the Fair Coin case. In the Fair Coin case, you have much more specific evidence. The evidence in the Mystery Coin example is too unspecific to warrant any assignment of a precise credence. It requires a different kind of uncertainty.

Because you have no information about the chance in the Mystery Coin case, this proponent of imprecise credences claims, you should have an imprecise credence that contains all of the probabilities that could be equal to the objective chance of the coin's landing heads, given your evidence. On this view, any precise credence would amount to taking a definite stance when the evidence doesn't justify a definite stance. It would mean adopting an attitude that was somehow more informative than what the evidence warrants.

The view that imprecise credences can be epistemically required is widespread:

[E]ven if men have, at least to a good degree of approximation, the abilities bayesians attribute to them, there are many situations where, in my opinion, rational men *ought not* to have precise utility functions and precise probability judgments. (Levi 1974, 394–5)

If there is little evidence concerning Ω then beliefs about Ω should be indeterminate, and probability models imprecise, to reflect the lack of information. (Walley 1991)

Precise credences . . . always commit a believer to extremely definite beliefs about repeated events and very specific inductive policies, even when the evidence comes nowhere close to warranting such beliefs and policies. (Joyce 2010, 285)

If you regard the chance function as indeterminate regarding X , it would be odd, and arguably irrational, for your credence to be any sharper . . . How would you defend that assignment? You could say "I don't *have* to defend it—it just happens to be my credence."

52 | Jennifer Rose Carr

But that seems about as unprincipled as looking at your sole source of information about the time, your digital clock, which tells that the time rounded off to the nearest minute is 4:03—and yet believing that the time is in fact 4:03 and 36 seconds. Granted, you may just happen to believe that; the point is that you have no business doing so.

(Hájek and Smithson 2012, 38–9)

Now, it's clear that there is no barrier to distinguishing the states of uncertainty in the Fair Coin and Mystery Coin cases within the precise credence model. Whether an agent is certain about the objective chance of heads or deeply uncertain is not obviously the kind of thing that can be read locally off of the agent's credence in the proposition that the coin will land heads. There are other relevant global features of the agent's total doxastic state: for example, the agent's credences in propositions explicitly about the objective chance of heads (e.g. low credence that $Ch(\text{heads}) = 0.5$).⁴⁰ In the Fair Coin case, the rational agent with precise credences adopts a credence function which assigns $Pr(\text{heads}) = 0.5$ and $Pr(Ch(\text{heads}) = 0.5) = 1$. In the Mystery Coin case, she adopts a credence function which assigns $Pr(\text{heads}) = 0.5$ and $Pr(Ch(\text{heads}) = 0.5) \approx 0$. These are different total doxastic states; so the precise credence model can distinguish different forms of uncertainty appropriate to the two evidential situations.

In general, for any specification of what information precise credences are meant to be inappropriately committed to, it will always be possible to characterize a precise credence function that is entirely noncommittal about (i.e., that has middling credence in) that information. The information might be information about objective probabilities, or evidential probabilities, etc. But even if that weren't the case—even if no such proposition could be specified in terms of, e.g., possible worlds propositions—we're still not forced to represent the relevant imprecision (suspension) in the attitude. We might use non-possible-worlds propositions that the agent could have precise credence in. (Propositions as sets of $\langle w, Pr \rangle$ pairs, for example.) For that matter, we could represent the same state of uncertainty using binary belief.

Where an agent is uncertain about some probabilistic features of a proposition p (e.g., its objective chance), there is a range of probabilities of p that the agent treats as open: call this set S . The imprecise credence model uses a measurement system that treats S as a feature of the attitude toward the content p . The precise credence model uses a measurement system that treats all non-subjective probabilities in S as components of contents: propositions about the probability of p . Only subjective probabilities figure into the attitude.

⁴⁰ The degree of resilience in the agent's credence in heads (that is, how much her credence in heads would be affected by the introduction of new evidence) is also relevant. See Skyrms (1977).

Relativism about the content/attitude distinction predicts: this disagreement has only to do with the assessor's choice of measurement system, and not with the agent's psychological or decision-theoretic commitments.

In section 2 I argued that the measure theory of mental content generated three forms of arbitrariness in the assignment of contents. Each of these forms of arbitrariness potentially affects the modeling of subjective probability:

1. There is arbitrariness in the assignment of indices to physical properties.

This is most obvious in the scale for subjective probability. It is conventional to treat 0 as the lowest subjective probability and 1 as the highest; a scale from -22 to 84, where -22 represents maximal certainty in a proposition and 84 represents maximal rejection, would have represented cardinal uncertainty just as well (though it would have made for uglier math). The measure theorist will say that there are also multiple measurement systems for *contents* that plausibly do the job equally well: for example, different conceptions of propositions.

2. There is arbitrariness in the choice of which physical properties to measure.

The measure theorist should say: in content measurements, we may include whatever cluster of dispositions form the ground for subjective probabilities as well as commitment to other forms of probability, like the proponent of probabilistic binary belief. Or we might include other forms of probability (in particular objective probability) within the contents, but exclude subjective probability, treating it as a different measurement, like the proponent of precise credences. Or we might exclude all psychological dispositions related to probability, treating them all as figuring into the attitude, like the proponent of imprecise credences.

3. There is arbitrariness in the individuation of measurements.

In the mystery coin case, the imprecise credence model sees one attitude toward one content where the precise credence model sees multiple attitudes toward multiple contents (credence in p versus credence in propositions about the chance of p). The measure theory predicts that these are equally accurate: again, compare measuring velocity to measuring speed and separately measuring direction.

4. Ramifications for Epistemology

Section 2 argued that if the measure theory of mental content is true, then the content/attitude distinction is measurement system relative (premise 1). Section 3 showed that if the content/attitude distinction is measurement system relative, then there's no psychological difference between having probabilistic binary beliefs, having precise credences, or having imprecise credences (premise 2). Now we can finally turn to the final premise in the central argument of this chapter:

- (3) If there's no psychological difference between having probabilistic binary beliefs, having precise credences, and having imprecise credences, then none of these is rationally impermissible.

From these three premises, it follows that if the measure theory of mental content is true, then there can be no rational requirement against probabilistic binary beliefs, or precise credences, or imprecise credences.

I've argued that the measure theory is committed to the view that the difference between these models boils down to where they encode subjective and objective probabilities: as figuring into the measurement that receives the title "content" or as figuring into other measurements of the attitude.

This conclusion, if correct, has normative implications. I claim the following: whether an agent is rational does not depend on the specific system of measurement that the assessor uses to characterize the agent's doxastic states.

Objection. What you've argued is compatible with the claim that precise or imprecise credences can be rationally impermissible on the measure theory. It's just that this impermissibility is relative to a model or measurement system. For example, *relative to the imprecise credence model*, it's rationally impermissible to have a precise credence in heads in the Mystery Coin case. Moreover, all rationality and irrationality attributions have determinate truth conditions partly in virtue of something other than the psychological properties of the addressee. The conventions adopted by the attributor affect whether the attribution is true, false, or neither. The attribution will have a specific language, and a specific measurement system. Insofar as there are any mental states in any circumstances, they are irrational only relative to a measurement system.

Reply. Whether an agent has precise credences, imprecise credences, or probabilistic binary beliefs depends partly on the assessor's choice of measurement system. Whether the agent is *rational* does not.⁴¹

A somewhat more committal version of this reply will pair the measure theory of mental content with evidentialism.⁴² On this view, whether an agent is epistemically rational is a function of her doxastic states and her evidence. Whether she is pragmatically rational is a function of her doxastic and orrectic states and her evidence. All three factors are entirely independent of the theorist's choice of representational conventions. And so whether the agent is rational or irrational is entirely independent of the theorist's representational conventions.

⁴¹ It's consistent with this thesis that rationality attributions are context- or assessment-sensitive in some other way: for example, it may be sensitive to norms of rationality.

⁴² On my interpretation, evidentialism of this sort has been endorsed by major figures on different sides of the debate: for example, Joyce (2010), who argues that precise credences are sometimes rationally impermissible, and Elga (2010), who argues that imprecise credences are always rationally impermissible.

If the measure theory is correct, then whether an agent is correctly represented as having precise credences, imprecise credences, or binary beliefs about subjective probability is a matter of the assessor's system of measurement. So, I conclude, it cannot be a requirement of rationality that an agent have imprecise credences, or precise credences, or binary beliefs about subjective probability. It is no rational failing of the agent that some theorist uses one measurement system or another to represent her mental states.

The main argument of this chapter is intentionally left with a conditional conclusion. The measure theory does face significant challenges. It's not clear, for example, that the measure theory is consistent with every theory of propositions or contents⁴³ or with every theory of the functional role of doxastic states. Nevertheless, given how perennially attractive the view has proven to be, it's worth exploring its consequences.

The theory of intentionality and the theory of rational belief are typically treated as orthogonal. It's clear that we have contentful mental states. The question in philosophy of mind of how exactly doxastic states get their content seems independent of the epistemological question of which doxastic states are rational under which circumstances. But some theories of intentionality have significant consequences in epistemology. If normative theories in epistemology are stated in terms of requirements for types of psychological state, we had better have a clear sense of what these psychological states amount to, how they differ from impermissible alternatives, and whether the differences are mere artifacts of our systems for representing the mind.⁴⁴

References

- Carr, Jennifer (2013). "Imprecise Evidence without Imprecise Credences." (manuscript.)
- Churchland, Paul M. (1979). *Scientific Realism and the Plasticity of Mind*. Cambridge University Press.
- Davidson, Donald (1973). "Radical Interpretation." *Dialectica*, 27(1): pp. 314–28.
- Davidson, Donald (1989). "What is Present to the Mind?" In Grazer Philosophische Studien. Amsterdam: Rodopi, pp. 197–213.
- de Finetti, Bruno (1937). "La Prévision: Ses Lois Logiques, Ses Sources Subjectives." *Annales de l'Institut Henri Poincaré*, 17: pp. 1–68.
- Dennett, Daniel C. (1987). *The Intentional Stance*. Cambridge, MA: MIT Press.

⁴³ In particular, there may be tensions between the measure theory and possible worlds propositions.

⁴⁴ Many thanks to Seamus Bradley, Annalisa Coliva, Samuel Fletcher, Rachel Goodman, Hanti Lin, Daniel Malinsky, Conor Mayo-Wilson, Julia Staffel, Michael Titelbaum, Greg Wheeler, Robbie Williams, and audiences at the Centre for Metaphysics and Mind at the University of Leeds (2015), the European Philosophy of Science Association (2015), and the Pacific APA (2017). Special thanks to Jonathan Livengood and Eric Pacuit. The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement n. 312938.

- Dretske, Fred (1981). *Knowledge and the Flow of Information*. Cambridge, MA: MIT Press.
- Dretske, Fred (1988). *Explaining Behavior: Reasons in a World of Causes*. Cambridge, MA: MIT Press.
- Egan, Andy (2007). "Epistemic Modals, Relativism and Assertion." *Philosophical Studies*, 133(1): pp. 1–22.
- Elga, Adam (2010). "Subjective Probabilities Should Be Sharp." *Philosophers' Imprint*, 10(05).
- Field, Hartry (1980). *Science Without Numbers*. Princeton, NJ: Princeton University Press.
- Fine, T. L. (1973). *Theories of Probability*. New York: Academic Press.
- Fitelson, Branden (2016). "Coherence." (manuscript).
- Fodor, Jerry A. (1987). *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*. Cambridge, MA: MIT Press.
- Frege, Gottlob (1892). "Über Sinn und Bedeutung." *Zeitschrift für Philosophie und philosophische Kritik*, 100: pp. 25–50.
- Gibbard, Allan (1990). *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Cambridge, MA: Harvard University Press.
- Gibbard, Allan (2003). *Thinking How to Live*. Cambridge, MA: Harvard University Press.
- Hájek, Alan and Michael Smithson (2012). "Rationality and Indeterminate Probabilities." *Synthese*, 187(1): pp. 33–48.
- Harman, Gilbert (1986). *Change in View*. Cambridge, MA: MIT Press.
- Harman, Gilbert and Judith Jarvis Thompson (1996). *Moral Relativism and Moral Objectivity*. Oxford: Blackwell.
- Holton, Richard (2014). "Intention as a Model for Belief." In M. Vargas and G. Yaffe (eds) *Rational and Social Agency: Essays on the Philosophy of Michael Bratman*. Oxford: Oxford University Press, pp. 12–37.
- Joyce, James M. (2010). "A Defense of Imprecise Credences in Inference and Decision Making." *Philosophical Perspectives*, 24(1): pp. 281–323.
- Keynes, John Maynard (1921). *A Treatise on Probability*. New York: Dover Publications.
- Koopman, B. O. (1940). "The axioms and algebra of intuitive probability." *Annals of Mathematics*, 2(41): pp. 269–92.
- Krantz, D. H., Suppes P., Luce, R. D., and Tversky, A. (1971). *Foundations of Measurement, Volume I*. New York: Academic Press.
- Kriegel, Uriah (2013). *Phenomenal Intentionality*. Oxford: Oxford University Press.
- Kyburg, Henry (1983). *Epistemology and Inference*. Minneapolis, MN: University of Minnesota Press.
- Levi, Isaac (1974). "On Indeterminate Probabilities." *Journal of Philosophy*, 71(13): pp. 391–418.
- Lewis, David (1974). "Radical Interpretation." *Synthese*, 23: pp. 331–4.
- MacFarlane, John (2014). *Assessment Sensitivity: Relative Truth and its Applications*. Oxford: Oxford University Press.
- Matthews, Robert J. (1994). "The Measure of Mind." *Mind*, 103(410): pp. 131–46.
- Millikan, Ruth G. (1989a). "Biosemantics." *Journal of Philosophy*, 86(July): pp. 281–97.
- Millikan, Ruth G. (1989b). "In Defense of Proper Functions." *Philosophy of Science*, 56 (June): pp. 288–302.
- Moss, Sarah (2015). "Credal Dilemmas." *Noûs*, 48(3): pp. 665–83.

Subjective Probability and Content/Attitude | 57

- Pedersen, Arthur Paul and Gregory Wheeler (2014). "Demystifying Dilation." *Erkenntnis*, 79(6): pp. 1305–42.
- Russell, Bertrand (1903). *Principles of Mathematics*. Cambridge: Cambridge University Press.
- Savage, Leonard (1954). *The Foundations of Statistics*. New York: John Wiley and Sons.
- Savage, Leonard (1972). *The Foundations of Statistics*, 2nd revised edition (first published 1954). New York: Dover.
- Schroeder, Mark (2009). "Being For: Evaluating the Semantic Program of Expressivism." *Analysis*, 70(1): pp. 101–4.
- Seidenfeld, Theodore (2004). "A Contrast Between Two Decision Rules for Use with (Convex) Sets of Probabilities: Γ -Maximin Versus E-Admissibility." *Synthese*, 140: pp. 69–88.
- Sepielli, Andrew (2012). "Subjective Normativity and Action Guidance." In M. Timmons (ed.) *Oxford Studies in Normative Ethics*, Vol. II. Oxford: Oxford University Press, pp. 45–73.
- Skyrms, Bryan (1977). "Resiliency, Propensities, and Causal Necessity." *Journal of Philosophy*, 74: pp. 704–13.
- Soames, Scott (2014). *Propositions as Cognitive Event Types*. Oxford: Oxford University Press.
- Stalnaker, Robert (1984). *Inquiry*. Cambridge, MA: MIT Press.
- Stampe, Dennis W. (1977). "Towards a Causal Theory of Linguistic Representation." *Midwest Studies in Philosophy*, 2(1): pp. 42–63.
- Sturgeon, Scott (2008). "Reason and the Grain of Belief." *Noûs*, 42(1): pp. 139–65.
- Suppes, Patrick and Zinnes, Joseph L., (1963). "Basic Measurement Theory," in R. D. Luce, R. R. Bush, and E. Galanter (eds), *Handbook of Mathematical Psychology*, Vol. 1. New York: Wiley, pp. 1–76.
- Tye, Michael (2000). *Consciousness, Color, and Content*. Cambridge, MA: MIT Press.
- Walley, Peter (1991). *Statistical Reasoning with Imprecise Probabilities*. London: Chapman & Hall.
- White, Roger (2009). "Evidential Symmetry and Mushy Credence." In Tamar Szabo Gendler and John Hawthorne (eds), *Oxford Studies in Epistemology*. New York: Oxford University Press, pp. 161–86.
- Williams, J. Robert G. (manuscript). "Reductivist Theories of Intentionality and the Representation of Options."
- Williamson, Timothy (2009). "The Philosophy of Philosophy." *Analysis*, 69(1): pp. 99–100.