

Philosophy of Science Association

Representation Theorems and Realism about Degrees of Belief

Author(s): Lyle Zynda

Source: *Philosophy of Science*, Vol. 67, No. 1 (Mar., 2000), pp. 45-69

Published by: The University of Chicago Press on behalf of the Philosophy of Science Association

Stable URL: <http://www.jstor.org/stable/188613>

Accessed: 12/03/2010 02:46

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=ucpress>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



The University of Chicago Press and Philosophy of Science Association are collaborating with JSTOR to digitize, preserve and extend access to *Philosophy of Science*.

Representation Theorems and Realism About Degrees of Belief *

Lyle Zynda^{†‡}

Indiana University South Bend

The representation theorems of expected utility theory show that having certain types of preferences is both necessary and sufficient for being representable as having subjective probabilities. However, unless the expected utility framework is simply assumed, such preferences are also consistent with being representable as having degrees of belief that do not obey the laws of probability. This fact shows that being representable as having subjective probabilities is not necessarily the same as having subjective probabilities. Probabilism can be defended on the basis of the representation theorems only if attributions of degrees of belief are understood either antirealistically or purely qualitatively, or if the representation theorems are supplemented by arguments based on other considerations (simplicity, consilience, and so on) that single out the representation of a person as having subjective probabilities as the only true representation of the mental state of any person whose preferences conform to the axioms of expected utility theory.

1. Introduction. Probabilists maintain that belief comes in degrees, and that (ideally) rational people have degrees of belief that conform to the laws of probability and so can be referred to properly as *subjective probabilities*. These two theses constitute the fundamental synchronic claims of probabilism. They are typically supplemented by various diachronic

*Received November 1998; revised September 1999.

[†]Send requests for reprints to the author, Department of Philosophy, Indiana University South Bend, 1700 Mishawaka Avenue, P.O. Box 7111, South Bend, IN, 46634.

[‡]I wish to thank Alan Hájek, Michael Kinyon, J. Wesley Robbins, and the participants at the Conference on Methods (New School, New York City, May 15, 1999) for helpful discussion of this essay. Thanks also to one of the anonymous referees, whose insightful comments prompted me to clarify several points.

Philosophy of Science, 67 (March 2000) pp. 45–69. 0031-8248/2000/6701-0003\$2.00
Copyright 2000 by the Philosophy of Science Association. All rights reserved.

theses detailing how subjective probabilities should be updated as new information is acquired.¹

Over the years, a large variety of arguments have been put forward for probabilism.² One long-popular class of arguments—Dutch book arguments—have increasingly fallen into disfavor after enduring a great deal of criticism.³ Many probabilists today regard Dutch book arguments only as useful illustrations or dramatizations of deeper truths about rational preference, truths stated more precisely by the representation theorems of axiomatic expected utility theory, upon which the case for probabilism is supposed to be properly grounded.

2. Representation Theorems as the Basis for an Argument for Probabilism.

In broad terms, a representation theorem shows that if a preference ordering has certain formal properties, then it is possible to define a certain type of real-valued function R that reproduces or “mirrors” that ordering, in the following sense.

Representation: $A > B \Leftrightarrow R(A) > R(B)$

Here A and B are the objects of preference—e.g., acts, propositions, probability distributions on an outcome space (lottery acts), etc., depending on the theory—and “ $A > B$ ” means “ A is strictly preferred to B .” Following Savage (1954), we will think here of A and B as *acts*, conceived of as functions from possible states of the world s_i to outcomes o_i . If the representing function R can be expressed in the form $\sum_i p(s_i)u(o_i)$, where p is a probability function defined on the Boolean algebra generated by $\{s_i\}$ (the algebra consisting of events E , which are sets of states of the world) and u is a utility function defined on $\{o_i\}$, it is an expected utility function. There are several representation theorems for expected utility, each based on a slightly different set of axioms for preference, as well as a number of alternative non-expected utility representation theorems.⁴

An argument for probabilism based on a representation theorem goes as follows. First, it is argued that certain axioms for preference in expected utility theory describe properties that hold of all rational preference. These

1. An excellent discussion of the various methods for updating subjective probabilities can be found in Jeffrey 1988.
2. Howson and Urbach (1989) and Earman (1992) survey and discuss many of these arguments.
3. Dutch book arguments were first stated explicitly by Ramsey (1926) and de Finetti (1937). Recent examples of probabilists who have questioned traditional Dutch book arguments include Maher (1993) and Kaplan (1996).
4. See Fishburn 1981, 1982 for a survey of expected utility theories, and Fishburn 1988 for a survey of alternatives to expected utility theory.

axioms vary somewhat from theory to theory, but all expected utility theories require certain things, such as that preferences be asymmetric, transitive, and that they have a property commonly referred to as independence.⁵ Other axioms (such as the various continuity or convexity axioms) are typically added to ensure the existence of a real-valued, continuous representation. These latter axioms are not usually defended as rationality axioms per se, but as technical necessities. They require that a person's preferences be extremely and arguably unrealistically rich and extensive. Probabilists often deal with these technical axioms by arguing that rationality requires only that one's preferences be embeddable in a continuous (or convex) set of preferences of the sort assumed by the representation theorems.⁶ Preferences that violate the axioms aptly regarded as rationality conditions (such as asymmetry, transitivity, and independence) are not so embeddable. The second step of the argument is an appeal to the representation theorems, which as purely mathematical theorems are beyond question. This stage of the argument involves the descriptive assumption, which we will examine in detail in what follows, that a person whose preferences conform to the specific rationality axioms assumed by the representation theorem really has values that can be measured by the utility functions defined by the representation, and degrees of belief similarly defined that conform to the laws of probability, which are as follows:

Nonnegativity. $p(E) \geq 0$ for all events E .

Normality. $p(S) = 1$, where S is the necessary event (the set of all possible states of the world).

Additivity. $p(E_1 \cup E_2) = p(E_1) + p(E_2)$ if E_1 and E_2 are mutually exclusive events ($E_1 \cap E_2 = \emptyset$).

There are two preliminary comments I would like to make about this framework before we begin our main discussion. First, even when a person's preferences conform to all the axioms of expected utility theory, including the continuity (or convexity) axioms, there will not in general be only one probability-utility function pair that represents the person's preferences. For example, in Savage's (1954) system, the probability function p is unique (given the conventional choice of a 0-to-1 scale for probabilities) but the utility function u is "unique" only up to a positive linear transformation. The probability function combined with *any* of these utility functions according to the formula $\sum_i p(s_i)u(o_i)$ will produce a function that represents (mirrors the order of) the person's preferences; thus, any two of these probability-utility function pairs, when combined according

5. Since the work of Allais (1952), independence axioms have been a primary target for those wishing to criticize the normative appropriateness of expected utility theory.

6. See, e.g., Skyrms 1984, 1987.

to the expected utility formula, produce *ordinally equivalent* representations of those preferences. The different utility functions can be regarded as distinct but equally valid scales for measuring subjective value, just as the Fahrenheit, Celsius, and Kelvin scales provide equally valid ways of measuring temperature.⁷ In Jeffrey's (1983) system, the probability function is forced to be unique (even given the choice of a 0-to-1 scale) only if the utilities are unbounded above and below. In other cases, a number of probability functions will represent the person's degrees of belief equally well, even though the person has an extremely rich and extensive set of preferences (defined, in Jeffrey's system, on an "atomless" Boolean algebra of propositions from which the impossible proposition has been removed).⁸ Moreover, if a person's preferences are not as extensive as the representation theorems assume, but can be embedded in (extended to) a structure that conforms to the axioms assumed by those theorems, there will be a correspondingly larger class of probability-utility function pairs that can represent (reproduce the order of) the person's preferences. Sec-

7. Subjective utility functions as defined by the representation theorems and the temperature scales in common use (Kelvin, Celsius, Fahrenheit, etc.) are similar in that each forms a class with the property that every member of the class is a linear transformation of any another member. This feature is part of what makes utility a *cardinal* measure of degree of subjective value, as opposed to a mere *ordinal* representation of a value ranking. An ordinal representation O simply reproduces an ordering, so that one cannot interpret the fact that, say, $O(x)$ is twice as great as $O(y)$ as meaning that x has twice the amount of some property P than y . The higher number represents a higher ranking in the ordering, and nothing more. A cardinal measure, by contrast, represents degrees of strength or intensity in addition to order. One invariant property of the different temperature scales commonly used is ratios of differences between temperatures $[(t_1 - t_2)/(t_3 - t_4)]$, which remain unchanged with a change of scale. Thus, while we cannot sensibly say that 200° F is "twice as hot" as 100° F, since the corresponding relationship does not hold if we switch to the Kelvin scale, we can say that the difference in temperature between 200° F and 100° F is twice as great as the difference between 100° F and 50° F, since this ratio will remain the same if we switch to the Celsius or Kelvin scales. That ratios of temperature differences are invariant is implied by the fact that the various temperature scales transform linearly into one another. Hence, (cardinal) utilities also have this property.

8. It is worth mentioning that in Jeffrey's system, utilities (which he calls desirabilities) are related to each other by fractional linear transformations. Specifically, where des and DES are two desirability functions that represent a person's values, then $DES(X) = [a des(X) + b]/[c des(X) + d]$, for all X and some a, b, c, d such that $ad - bc > 0$, $c des(X) + d > 0$, and $c des(T) + d = 1$, T being the necessary proposition. When desirabilities are unbounded both above and below, this forces $c = 0$ and $d = 1$, so that the fractional linear transformations reduce to simple linear transformations. Similarly, for all propositions X , if $PROB$ and $prob$ are two probability functions that represent the person's degrees of belief, then $PROB(X) = prob(X)[c des(X) - d]$, implying that $PROB(X) = prob(X)$ (the probability function is unique) only when c can only be 0, which occurs when and only when desirabilities are unbounded both above and below. See Jeffrey 1983, Ch. 6, for further discussion.

ond, it is important to emphasize that it is not just the probability-utility function pair that represents the person's preferences: it is the pair *in conjunction with* a method (mathematical expectation, as embodied in the formula $\sum_i p(s_i)u(o_i)$) for combining the elements of the pair into a representing function. Thus, a representation theorem defines *three* things that let us "represent" someone's preferences: (1) a probability function or class of such functions, (2) a corresponding class of utility functions (typically unique up to a positive linear transformation), and (3) a method, mathematical expectation, for combining (1) and (2) to mirror (reproduce the order of) the preference ranking.⁹

3. Representation Theorems and Realism about Degrees of Belief. An important descriptive question arises whether *being representable as having* certain degrees of belief, as described by a probability function p , is sufficient for *really having* those degrees of belief. One thing that is clear from the above discussion is that it is not, at least not in the simple formulation just stated. For if someone's degrees of belief can be represented by more than one probability function—which would occur if that person's preferences were not as extensive as the representation theorems assume, or if (in Jeffrey's system) his or her utilities were bounded either above or below—it would be wrong to say that any specific one of those functions describes that person's "real" degrees of belief. In such cases, most probabilists would grant that the person should be regarded as having *vague* or *indeterminate* (interval-valued) degrees of belief. Van Fraassen (1984, 1989) refers to the class of probability functions that are consistent with all of someone's probabilistic judgments as that person's *representor*. What is true of the person's opinion is what *all* of the probability functions in the person's representor have in common.¹⁰ Epistemic rationality on van Fraassen's approach requires only that a person's representor be non-empty. Probabilists who prefer a more pragmatic, explicitly decision-

9. I should note that in this essay I am assuming (as do the representation theorems) that the preference ordering is known. A substantive interpretive question exists concerning what facts determine the "right" preference orderings to attribute to a person. In particular, Hurley (1989) and Broome (1991) consider the question of whether purely formal conditions on preference such as transitivity and independence could by themselves constrain preferences at all, either descriptively or normatively, given that one can seemingly explain away any apparent violation of these conditions by redescribing the objects of preference in a more fine-grained way. Both Hurley and Broome argue that formal conditions are empty in the absence of substantive constraints that require agents to be indifferent between certain distinct options. Not any difference ought to make a difference to one's preferences.

10. There are a few exceptions to this—e.g., all of the probability functions in someone's representor are precise, but the person's opinion will not be precise if there is more than one function in the representor.

theoretic approach to opinion can also allow for vague degrees of belief. If there is more than one probability function in the class of probability-utility function pairs, each of which represents a person's preferences, then in analogy with van Fraassen's notion of a representor, what should be seriously and realistically attributed to the person's opinions and values will be what all such probability-utility pairs that represent the person's preferences have in common, and no more. This sort of limited indeterminacy is perfectly consistent with the spirit of the representation theorems, most of which, as we have seen, define utility only up to a positive linear transformation, anyway, and some of which (Jeffrey's theory) do not define a unique probability function even for certain highly idealized, rich preference structures.¹¹

It is evident from this that the problem some nonprobabilists have with the notion of subjective probabilities and utilities—namely, that they find it difficult to conceive seriously and realistically of people as having “numbers in the head” and as somehow (unconsciously?) calculating with those numbers when reasoning—is somewhat misplaced. The representational approach outlined above makes it clear that probabilists needn't be committed to a naïve “numbers in the head” account of opinion and reasoning. Probabilists can say that what is *literally true* of a person's opinion—what should be understood seriously and realistically in our attributions of degrees of belief to that person—are not the particular numbers, but the properties that are common to all probability functions that appear in at least one expected utility representation of his or her preferences. This would presumably include the properties described in the axioms of probability (nonnegativity, normality, additivity), the property of one's (vague) degrees of belief for propositions or events being certain intervals¹² (which for some propositions or events could be a closed interval $[x, x]$, in which case probabilities for those propositions or events would be

11. I should note a difference between the type of “indeterminacy” in probabilities allowed by Jeffrey's theory and the indeterminacy due to vague probabilities in other contexts. In Jeffrey's theory, the admissible probability functions can be transformed (with the desirability functions) into one another without making a difference to the person's (fully defined) preferences, so they are best considered different probability scales. When a person's probabilities are vague in a framework like Savage's, however, the case is different: the person's preferences cannot be fully defined, else the probability function would be uniquely determined. (My thanks to Isaac Levi for pointing this out to me.)

12. This formulation means that the person's vague subjective probability for a proposition or event is not really anywhere within the assigned interval, but is to be identified with that interval or understood as being “spread across” it. Thus, it does not mean that the person's “real” subjective probability for a proposition or event is sharp, and it is simply “unknown” where in a certain interval it lies.

sharp), and the particular ordinal relationships between propositions or events ranked as more or less probable.

To sum up, representation theorems in expected utility theory, properly understood, show the following with respect to degrees of belief.

Representability. If a person's preferences obey the axioms of expected utility theory, then he or she can be represented as having degrees of belief that obey the laws of the probability calculus.

If we make the following two assumptions

The Rationality Condition. The axioms of expected utility theory are the axioms of rational preference.

The Reality Condition. If a person can be represented as having degrees of belief that obey the probability calculus, then the person really has degrees of belief (possibly vague, if there is more than one such representation) that obey the laws of the probability calculus.

the statement below follows.

Thesis 1: If a person's preferences obey the axioms of rational preference, then the person has (possibly vague) degrees of belief that obey the laws of the probability calculus.

This logically implies (via contraposition) the following thesis.

Thesis 2: If a person does not have (possibly vague) degrees of belief that obey the laws of the probability calculus, then that person violates at least one of the axioms of rational preference.

Thus, granting the Rationality and Reality Conditions, the representation theorems can be used to provide an argument for probabilism: having degrees of belief that conform to the laws of probability is required if one is to have rational preferences. This means that progress toward defending probabilism can be made by defending the axioms of expected utility theory as axioms of rational preference.¹³ In what follows, however, I will concentrate not on this, the Rationality Condition, but on the Reality Condition.

4. The Problem of Alternative Representations of Rational Preferences. Let us suppose we have two friends, Leonard and Maurice. Leonard claims to have subjective probabilities. His preferences obey the axioms of expected utility theory, and so he claims (in line with the Reality Condition

13. See Maher 1993 for an excellent recent defense of probabilism along these lines. Also of interest is Broome 1991.

stated above) that he has degrees of belief $p(s_i)$ that obey the laws of the probability calculus, and a utility function $u(o_i)$ that he implicitly uses along with his subjective probabilities to determine the expected utility $EU(A) = \sum_i p(s_i)u(o_i)$ of any act A . Maurice, by contrast, claims that he does not have degrees of belief that obey the laws of probability. His degrees of belief, he claims, are defined according to what he calls a 1-to-10 *believability ranking* for each possible state of the world s_i , which we will designate as $b(s_i)$. The axioms governing Maurice's believability rankings are the following:

Minimality. $b(E) \geq 1$ for all events E .

Maximality. $b(S) = 10$, where S is the necessary event (the set of all possible states of the world).

Subadditivity. $b(E_1 \cup \dots \cup E_n) = [b(E_1) + \dots + b(E_n)] - (n - 1)$ if for every i, j in $\{1, \dots, n\}$ such that $i \neq j$, E_i and E_j are mutually exclusive events ($E_i \cap E_j = \emptyset$).

We can use these three axioms to calculate believability rankings for specific cases in the "classical" manner. For example, rolling a fair die can result in any one of six equally possible and mutually exclusive outcomes. Since at least one outcome must occur, the believability ranking for the proposition that at least one of the six outcomes 1, 2, 3, 4, 5, or 6 will result is 10. This means that each outcome has a believability ranking $x = 2.5$, since $b(\text{exactly one of the outcomes 1, 2, 3, 4, 5, or 6 will occur}) = 10 = 6x - 5$. Similarly, Maurice's likelihood that an even number will come up is $(2.5 + 2.5 + 2.5 - 2) = 5.5$ —the same as his likelihood that an odd number will come up. Finally, the likelihood that an even number or an odd number will come up is $(5.5 + 5.5 - 1) = 10$ —just as one would expect. Now, Maurice's degrees of belief (believability rankings) do not obey the laws of the probability calculus as stated earlier. Most notably, they violate the additivity axiom. For example, the probability that an even number comes up is just the probability that either a 2 or a 4 or a 6 comes up, but the believability rankings Maurice gives to each of these three events—2.5 apiece—do not add up to 5.5.

Now, is it the case that Maurice must violate one of the axioms of rational preference, as defined by expected utility theory? Maurice claims not, for, he says, he does not use his believability rankings to calculate and maximize the *expectation* of utility. Remember that we needed *three* elements to define a particular expected utility representation of a preference ranking: (1) a probability function, (2) a utility function, and (3) a method of combining the two (mathematical expectation) to mirror the person's preferences. Maurice points out that he can compensate for his "deficiency" with respect to (1) by choosing a different method of combining his degrees of belief with his utilities to produce his preferences.

According to Maurice, he maximizes a quantity he calls *valuation*, which compensates for his having non-additive degrees of belief.

Valuation. The valuation $V(A)$ of an act A is defined as $\sum_i (b(s_i)u(o_i) - u(o_i))$.

Now, to make the story short, it just so happens that Maurice has the same utility function as Leonard, and that his degree of belief function $b(s_i)$ is related linearly to Leonard's probability function $p(s_i)$ as follows: $b(s_i) = 9p(s_i) + 1$. (The reader can verify that this is consistent with the axioms for believability rankings just stated. In particular, the linear transformation of p embodied by b is not additive but is subadditive in the manner defined above.¹⁴) Given this, it follows that $V(A)$ is equal to $9\sum p(s_i)u(o_i)$. Therefore, since the following is true

$$9\sum p(s_i)u(o_i) > 9\sum p(s_j)u(o_j) \Leftrightarrow \sum p(s_i)u(o_i) > \sum p(s_j)u(o_j),$$

it follows that

$$V(A) > V(B) \Leftrightarrow EU(A) > EU(B)$$

and therefore that

$$A > B \Leftrightarrow V(A) > V(B).$$

In other words, Maurice and Leonard have the very same preference rankings, even though we can represent Maurice as having degrees of belief that do not obey the laws of probability!

Now the descriptive question we were concerned with earlier arises again: what fact of the matter, if any, determines that Maurice *really* does not have degrees of belief that obey the laws of probability and that Leonard *really* does? It is clear from the argument above that their preferences cannot by themselves do the job, for Leonard and Maurice have the same preferences. The fact that they can both be *represented as* having subjective probabilities (degrees of belief that obey the laws of probability) does not by itself settle the issue, either, for they can also both be represented as having degrees of belief that do not obey the laws of probability. This means that, as far as their state of opinion goes, representability is not enough to determine what their degrees of belief really are, or even what formal properties they have (e.g., whether they are additive or subadditive). To look at this situation from another (normative) angle, if we grant

14. Some readers may be inclined, based on an analogy with utility and temperature (see fn. 7), to argue that b is essentially a probability scale, since it is a linear transformation of one. I think that the violation of the additivity axiom shows that it is not, though b is certainly a function of a type that includes probability functions. I will discuss this point further in the next section.

that Maurice's self-description is correct and that he *does* have degrees of belief (namely, his believability rankings) that violate the laws of probability, we must also admit that this does him no practical harm, since his preferences (defined by valuation rather than the expectation of utility) are the same as Leonard's. This means that we cannot argue the normative point that Maurice *ought* to have degrees of belief that obey the laws of probability on the grounds that otherwise he will violate one of the axioms of rational preference! Thus, if we take Maurice's self-characterization at face value, we seem to have no ammunition to convince him to adopt degrees of belief that conform to the laws of probability. Maurice seems to be saying, with some plausibility, "The laws of rational preference do not dictate that my degrees of belief obey the laws of probability. If a person's degrees of belief violate one of the laws of probability as you have stated them, but the method by which his degrees of belief combine with utilities to form preferences differs appropriately from the standard way of doing things, he will not necessarily violate the laws of rational preference. Certainly I do not."

5. Realisms and Antirealisms about Degrees of Belief. The problem that the case of Maurice and Leonard presents to the probabilist who bases the case for probabilism on the representation theorems is similar in many ways to the problems that empirically equivalent theories pose for the scientific realist. Just as even the best evidence may logically underdetermine the choice between two scientific theories, a preference ordering may not be representable in only one way. This throws the Reality Condition into question, and with it the argument for probabilism based on the representation theorems outlined above. Opponents of probabilism may be inclined to interpret this situation as an indication that the representation theorems cannot provide a foundation for probabilism. For, it might be argued, the representation theorems can provide a reason to have degrees of belief that obey the laws of probability only if a person cannot both have degrees of belief that fail to conform to the laws of probability and preferences that conform to the axioms of rational preference. Maurice's self-ascribed believability rankings show that this is not the case. However, as a (moderate) probabilist myself,¹⁵ I think that it would be hasty to conclude that the representation theorems cannot play a role in justifying probabilism, though of course the argument for probabilism based on the representation theorems outlined above will require some supplementation or reinterpretation. In what follows, I will discuss the implications of the case of Maurice for the probabilist. The discussion will focus on the Reality Condition. How can we decide which representation of Maurice's

15. See, e.g., Zynda 1996.

degrees of belief is “true”? Does Maurice really have subjective probabilities, despite his claims, or does he not?

Before we can fruitfully address this question, it is important to clarify the ontological stances it is possible to take toward degrees of belief, utilities, and their relationships to one another and to preference. First, there is the position I will call *eliminativism*, which holds that entities such as degrees of belief and utilities are psychologically unreal and that we should not make use of them in our theories, whether they be in epistemology, philosophy of mind, or philosophy of science. The second view, which I will refer to as *antirealism*, holds that these entities can be used legitimately to devise a formal theory that validates a certain class of decisions as rational, but that the theoretical entities (degrees of belief and utilities) in such a theory should not or at least need not be taken seriously as referring to psychologically real states. It is enough for the purposes of the anti-realist that a rational person behaves *as if* he or she has subjective probabilities and utilities and chooses acts by maximizing expected utility. The third view, which I will refer to as *weak realism*, holds that degrees of belief and utilities can be truly and justifiably attributed to people as describing aspects of their psychological states, but that degrees of belief, utilities, and preferences should not be thought of as independently existing, interacting mental states. Rather, preferences are ontologically primary, and degrees of belief and utilities are defined by logical construction from preferences.¹⁶ Finally, there is the view I will call *strong realism*,

16. The positions most closely matching what I call weak realism are the various interpretive theories of subjective probability and utility, e.g., those espoused by Hurley (1989) and Maher (1993). The idea is that we attribute to a person the preferences that make sense of his or her behavior, in the sense of rationalizing it as well as can be done. As Maher (1993, 9) puts it, “an attribution of probabilities and utilities is correct just in case it is part of an overall interpretation of the person’s preferences that makes sufficiently good sense of them and better sense than any other competing interpretation does For present purposes, it will suffice to assert that if a person’s preferences all maximize expected utility relative to some p and u , then it provides a perfect interpretation of the person’s preferences to say that p and u are the person’s probability and utility functions.” Thus, conformance to the preference axioms of expected utility theory is asserted to be a (defeasible) constraint on the attribution of preferences. The attribution of subjective probabilities and utilities is “correct” (presumably, *true*) when the preferences attributed (under the best interpretation) conform to the expected utility axioms. However, these probabilities and utilities needn’t be thought of as explicitly or consciously represented, rather they are “essentially a device for interpreting a person’s preferences.”

It is important to note that Hurley and Maher, and other authors who consider similar issues, such as Broome (1991), are mostly concerned with non-expected utility theories and how expected utility theories can be defended against them. Their arguments focus on showing that a rational agent’s preferences ought to conform to the axioms of expected utility theory (what I call the Rationality Condition). They assume that, if the preferences do so, the attribution of subjective probabilities and utilities is

which conceives of degrees of belief, utilities, and preferences as independently existing, interacting mental states. A person forms degrees of belief and utilities with respect to certain states and outcomes, and then combines these to form preferences. On this view, expected utility theory is a hypothesis about how distinct psychological states interact (degrees of belief and utilities combine to form preferences), and it also provides normative constraints on how these entities ought to be structured (degrees of belief ought to be probabilities) and interact (one should use one's degrees of belief and utilities to define the expectation of utility, and act to maximize expected utility).¹⁷

That said, let us turn to a careful examination of the differences between Maurice and Leonard. If we are only concerned with Maurice and Leonard as “black boxes,” where all that is important is the overall function (map from acts A to numbers representing those acts' choiceworthiness) they compute in forming their preferences, then we have the situation illustrated in Figure 1. In this figure, Leonard is represented by the top box, and Maurice by the bottom box. Now, if this figure captures all that is important to Leonard's and Maurice's mental states, it would seem that there is nothing substantive that distinguishes them. To see this, note that $V(A)$ is a very simple order-preserving transformation of EU. In other words, V just is an expected utility function, since multiplying an expected utility function by a positive constant produces another expected utility function. One can express $V(A)$ in the form $\Sigma p(s_i)u^*(o_i)$, where $u^*(o_i) =$

then justified and “correct” (the Reality Condition). It is the latter assumption with which the present paper is concerned. Because these authors do not deal with the precise issue considered here, it is possible (even though it seems closest to their views as expressed above) they would not endorse what I call “weak realism.” Even if they would, it is possible they might accept only the view that degrees of belief, utilities, and preferences are not “independently existing, interacting mental states” without the further characterization (which I make) that degrees of belief and utilities are “defined by logical construction” from preferences. (My thanks to an anonymous referee for this latter point.)

17. Harman's “combinatorial explosion” argument (1986, Ch. 3) is aimed at a strongly realistic version of probabilism that assumes that subjective probabilities are explicitly represented and transformed (e.g., via conditioning and other updating rules) in a manner similar to our (external) explicit calculations with probabilities. His argument is that reasoning cannot be based on explicitly represented subjective probabilities, since that would be computationally intractable. He later (104–105) uses this conclusion as part of an argument that expected utility theory cannot provide an account of our (explicit) practical reasoning.

A great deal of work in artificial intelligence is concerned with showing how implementation of probabilistic and uncertain reasoning of various sorts is possible. See, e.g., Pearl 1988 and Pearl and Shafer 1990. Such models, however, do not always aim at psychological realism (i.e., they are not typically intended as models of actual psychological processes).




Figure 1.

$9u(o_i)$. Thus, if one is inclined to deny the independent reality of Maurice's degrees of belief and utilities or their status as interacting mental states—if all that matters is the overall function they compute (EU or V)—one might have a case for arguing that there is no substantial difference between him and Leonard.¹⁸

This is not, however, the only way of thinking about Leonard and Maurice. For example, if one thinks about Maurice's and Leonard's self-descriptions in a strongly realistic sense as describing features of certain interacting and independently existing cognitive states and processes, it is clear that their self-descriptions are distinct even though the overall functions they compute are not significantly different. On this way of thinking, it is important to open the black boxes and see what's inside, as illustrated in Figure 2. In the two diagrams in this figure, we think of Leonard (top) and Maurice (bottom) as each having three “modules,” a degree of belief module (m_1), a utility module (m_2), and a module for combining the two to form preferences (the module at the left, which we can designate m_3).¹⁹

Now, looking at things this way it is clear that valuation and expected utility are mathematically distinct, since the formulas in the two left boxes are. One could not for instance replace the believability ranking module in the lower diagram with a probability module, leaving the valuation module to the far left intact, without producing obviously irrational preferences.²⁰ Calculating valuation is different from calculating expected util-

18. In fact, since EU represents a certain preference ordering, then any other function that represents the same ordering will have to be an order-preserving transformation of EU. (My thanks to Michael Kinyon for this observation.)

19. The module designations (m_1 , m_2 , m_3) represent place in the modular design and the input-output connections associated with that place, not the particular internal function computed by the modules placed there.

20. For example, using the valuation formula to determine preferences with a probability function substituted for the believability ranking would result in preferences that violate dominance. Suppose that there are two equiprobable cases, p and not- p , so that each gets

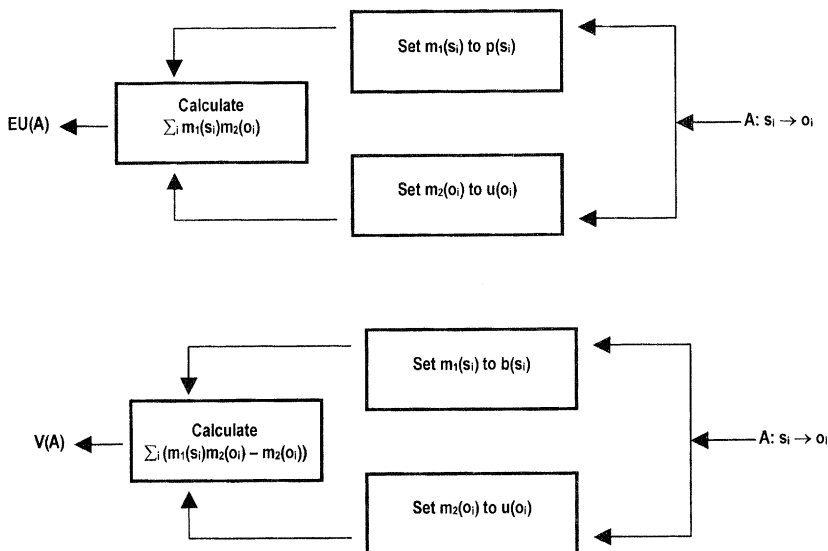


Figure 2.

ity, even if the results *with certain sorts of input* rank acts the same way. If someone were to program one computer with the top modular design (like Leonard), and another computer with the bottom modular design (like Maurice, as self-described), the two computers would have to have different programs, and would consequently go through a different sequence of internal states in making a decision. Therefore, for those inclined toward strong realism, there is a fact of the matter about whether Maurice is truly described by the top or bottom diagram of Figure 2 (or something else altogether). The strong realist has to decide between these two models based solely on considerations relevant to their *truth*. The problem for the strong realist is that Maurice's preferences and the Reality Condition cannot by themselves determine which is the correct choice.

Thus, the case of Leonard and Maurice does raise significant problems for the probabilist who (1) wants to be a strong realist about degrees of belief, utilities, and expected utility maximization as independently existing, in-

probability 1/2, and there are two acts, A and B. A has an outcome with utility 5 if p, and an outcome with utility 2 if not-p. B has an outcome with utility 4 if p, and an outcome with utility 1 if not-p. A is obviously the dominant act. However, using the probabilities just stated in the valuation formula results in $V(A) = ((5 \times 1/2) - 5) + ((2 \times 1/2) - 2) = -3.5$ and $V(B) = ((4 \times 1/2) - 4) + ((1 \times 1/2) - 1) = -2.5$, giving B the higher valuation. Using the believability ranking 5.5 for p and not-p, by contrast, preserves dominance.

teracting states and processes and (2) bases his or her defense of probabilism solely on the representation theorems. By contrast, an antirealist about these entities—one who thinks of degrees of belief, utilities, and expected utility maximization merely as convenient representations of the overall behavior of an unopened “black box”—could make a choice between modeling Maurice as having subjective probabilities (or believability rankings) on conventional or pragmatic grounds. For example, if the antirealist is an instrumentalist, and regards putative mental states such as degrees of belief as useful fictions only, then the choice between the two descriptions of Maurice’s degrees of belief can be justified on grounds of theoretical convenience, such as the adoption of a purely conventional definition of Maurice’s degrees of belief as what is common to all probability functions that can represent his preferences in the usual manner found in expected utility theory. Upon adopting this convention, we can attribute to Maurice the function $p(s_i)$ as found in the top box of Figure 2. On this view, attributions to an agent of specific degrees of belief, utilities, and their relationship to one another and to preference are relative to a theoretical framework, of which we can only use one at a time, on pain of internal inconsistency within our framework, but there is for an instrumentalist no serious ontological issue to settle in adopting one set of useful fictions over another. If the probabilist is an antirealist of the constructive sort, and thinks of probabilist models of opinion as possibly true or false descriptions of what’s “in the head” at some level of description but is agnostic on principle about which is true, the choice between which of the two frameworks to accept²¹ can be justified on pragmatic grounds, such as the usefulness of the framework to the theorizer, its scope, fruitfulness, simplicity, elegance, and formal tractability, and its similarity to theoretical frameworks already in use.²² For the antirealist, who is only concerned with “saving the phenomena” (here, the preference orderings regarded as rational are the “phenomena”), representability is enough; one needn’t worry about principles such as the Reality Condition.

Adopting an antirealistic stance toward degree of belief and utility would definitely allow a probabilist to continue to defend *using* the framework of expected utility and subjective probability and so avoid being pushed into the eliminativist camp. However, I suspect that many prob-

21. The term “accept” here should be understood in a manner similar to the notion of acceptance in van Fraassen’s version of scientific antirealism, known as constructive empiricism, in which one “accepts” a theory by believing it to be empirically adequate and committing oneself to reason and speak within the framework provided by the theory (van Fraassen 1980). Here mirroring a preference ranking plays the role of empirical adequacy.

22. Antirealists such as van Fraassen hold that these qualities are reasons to use a theory (hence the label “pragmatic”), but are evidentially irrelevant to a theory’s truth.

abilists wishing to remain realists of some sort about degrees of belief and utilities will find the antirealist “solution” to the problem of Maurice unsatisfactory. After all, both the eliminativist and the antirealist refuse to commit themselves to the literal truth of existence claims about degrees of belief and utilities. For the realistically inclined, a probabilist of the anti-realistic sort may talk like a probabilist but he or she is really a wolf (eliminativist) in sheep’s clothing. But what basis does a probabilist have to be a realist about these entities in light of Maurice’s alternative representation of his mental states, as illustrated in the bottom diagram of Figure 2?

6. Options for the Realist. The case of Maurice and Leonard highlights the tight connection between the descriptive and normative aspects of the concepts of degree of belief, utility, and preference. At issue between Leonard and Maurice is the normative question of whether a person can violate the laws of probability without violating the laws of rational preference. This in turn depends on the descriptive issue of exactly which degrees of belief, utilities, and method of combining the two can be justifiably attributed to them. The fact that the normative and descriptive aspects of this issue are so closely intertwined creates difficulties for the realist. Before we return to our discussion of the problems facing the strong realist, let us turn to weak realism, which we have not yet discussed, and see if the same problems affect this position as affected strong realism. The weak realist, you recall, holds that preferences are ontologically primary, and that degrees of belief and utilities are logical constructions from preferences. Thus, unlike the strong realist, the weak realist is essentially a reductionist about degrees of belief and utilities. According to weak realism, degrees of belief and utilities are real but have no existence independent from preferences. To use an analogy due to Daniel Dennett, for the weak realist, degrees of belief and utilities are “real” in the same sense that centers of gravity are real: centers of gravity are precisely defined by the mass distribution of a body, and claims about centers of gravity can be judged true or false based on information about ontologically prior entities (particle masses and positions) and well-defined rules (a mathematical formula for calculating what the center of gravity of an object is from these particle masses and positions). However, a center of gravity is not a further component or non-reducible property of the body. So, the weak realist’s claim is essentially that degrees of belief and utilities are *abstracta* (logical constructions from a person’s preferences) rather than *illata* (independently existing components of a person’s overall mental state or cognitive decision-making system).²³ If we adopt this view, then we seem-

23. For further discussion, see the essays in Dennett 1987, 1998. I do not claim here

ingly only have the task of giving a precise definition of terms such as “degree of belief” and “utility” by logical construction from the person’s preference ranking. This is seemingly less than what would be required of the strong realist. To this end, might not a weak realist simply take the view that Maurice and Leonard (and people in general) are “black boxes” as portrayed in Figure 1, since the ontologically basic items (preferences) are fully represented by the functions EU and V present in that figure (both of which are expected utility functions and so can be expressed in the form $\sum p(s_i)u(o_i)$), and simply *postulate* that if a person’s preferences obey the axioms of expected utility theory, then his or her degrees of belief are *by definition* the subjective probabilities (vague, if necessary) as defined by the standard representation theorems?

Unfortunately, if we do so, then we cannot give an independent defense of the standard definition of degree of belief (and of the normative requirement that degrees of belief be subjective probabilities) against Maurice’s objections, since we have simply by fiat eliminated his proposed alternative. Maurice can legitimately object to our doing so, since our postulation is not an innocent stipulative definition (as it would be if we defined a person’s “splork” toward a proposition as the probability provided by the standard representation theorem, in which case a person might have *both* a “splork” and a “believability ranking” as Maurice defines it), but a theoretical definition, in that postulating it *simply assumes* the standard method of defining degrees of belief within expected utility theory over Maurice’s proposed alternative. His proposal is that we decompose the function V (or EU) into two component functions b and u, combined according to his “valuation” formula $\sum_i (b(s_i)u(o_i) - u(o_i))$, and take b (which obeys his alternative axioms, including subadditivity) to represent his degrees of belief. It doesn’t address his proposal to say that we can express V (or EU) in the form $\sum p(s_i)u(o_i)$ and take p to measure his degrees of belief. That is a counterproposal, to be sure, but simply offering a counterproposal doesn’t by itself say *why* we shouldn’t do things Maurice’s way. Now, the concept of degree of belief is to a large extent a pre-theoretical notion, which for its initial intuitive appeal draws on the familiarity and usefulness of folk psychological categories such as belief and confidence. We can explicate this pre-theoretical notion by giving a precise definition within the framework of a formal theory, such as stan-

that Dennett’s much-discussed distinction between abstracta and illata is crystal clear, nor do I claim that it will necessarily be of comfort to all probabilists who wish to be realists about degrees of belief. The reason for this is that Dennett has been accused of being an instrumentalist about belief and desire, on the grounds that abstracta as he defines them would not be literally real, a charge Dennett has consistently denied. In my view (though I will not argue it here), the abstracta/illata distinction has some usefulness in deflecting the charge of instrumentalism.

standard expected utility theory, but the appropriateness of any such definition and the precise conception of degree of belief that it entails stands or falls with the theory. Hence, we cannot *defend* the standard approach by appealing to such a definition. Another reason we should not simply identify the concept of degree of belief with the specific definition in standard expected utility theory is that we can make sense of people having degrees of belief when their preferences do not obey the axioms of expected utility theory, and even sometimes give precise definitions of those degrees of belief.²⁴ If their degrees of belief are just *defined* to be what is common to all subjective probabilities that occur in some expected utility representation of their preferences, then such people would have no degrees of belief, since no expected utility representation of their preferences exists.²⁵ Therefore, since Maurice and Leonard are putting forward different theoretical developments of a pre-theoretical notion, the only way to decide between them is to compare their proposals directly. We cannot simply postulate that Maurice is wrong.

That said, a probabilist who wants to be a weak realist must either argue (1) that Maurice's definition of his degrees of belief is inferior to the standard one, or (2) that it is not significantly different from the standard one in those respects that are important to attributing degree of belief. Let us begin by examining strategy (1). Since the overall accounts that Maurice and Leonard offer are ordinally equivalent, no distinction can be made between the two approaches based on whether they correctly mirror their preference rankings. Thus, we must decide between the two proposed def-

24. See Fishburn 1988 for a survey of representation theorems for preferences that violate the axioms of expected utility theory. Useful discussion can also be found in §II of Gärdenfors and Sahlin 1988.

25. It might be objected that on the interpretive approach to preference and subjective probabilities, perfect fit to the expected utility axioms for preference is not required for the meaningful attribution of subjective probabilities; all that is needed is that the interpretation that provides the "best fit" (which might not fit *all* of a person's apparent preferences) attributes preferences that conform to those axioms. It is true that interpretive approaches do not require perfect fit: a single intransitivity in a person's preferences can be ignored if that person's preferences otherwise adhere to the axioms. However, I think it is an open question whether for some agents, axioms of preference from one of the non-expected utility theories will provide the "best fit." (This cannot be ruled out a priori; otherwise, the vast literature on non-expected utility theories would be meaningless, which it is not. There is room for reasonable disagreement on both the descriptive and normative appropriateness of expected utility theory.) Moreover, I want to leave open that the best interpretation will sometimes require one to attribute degrees of belief to an agent that violate the probability axioms. If this is possible, however, then there would have to be a way of specifying those degrees of belief that does not depend on the standard representation theorems of expected utility theory, since those theorems provide resources for specifying degrees of belief when and only when the assumed preference axioms apply.

initions based on their internal virtues. Moreover, since we are attempting to defend a form of realism, we must appeal only to those virtues that are arguably relevant to the truth²⁶ of the two theories (which according to the realist might include things such as simplicity, elegance, fruitfulness, scope, strength, consilience, and so on). I cannot of course address the truth-relevance of such virtues here (other than to remark that if the reader is skeptical about the truth-relevance of these “theoretical virtues,” he or she will have to be correspondingly skeptical about the defense of weak realism about degrees of belief developed here), but it is possible to discuss how such an argument could be developed. A weak realist might cite several considerations to justify a decision to accept the standard definition of degree of belief in expected utility theory over Maurice’s definition, such as consilience (probability theory is well-established in mathematics, statistics, and economics), or simplicity or “naturalness” (perhaps Maurice’s definition of degree of belief seems somewhat more complex or “contrived” than the standard one). The strong realist (who thinks of degrees of belief, utilities, and preferences as *illata*—independently existing, interacting mental states) could adopt similar strategies to supplement the argument from the representation theorems, arguing that the top diagram in Figure 2 is inherently superior (simpler, more natural, etc.) than the bottom diagram. In either case, doing so requires going beyond a simple appeal to Maurice’s preferences and the Reality Principle. It requires that the case for probabilism be supplemented by an argument showing that the standard definition of degree of belief is the simplest, most fruitful, etc., of all alternative definitions, including Maurice’s definition (believability rankings). Now, it seems to me that such a defense could possibly be successful in the case of the particular alternative put forward by Maurice. Let us consider only one of the theoretical virtues, simplicity. Note first that there are extra terms in the definitions of valuation and Maurice’s subadditivity axiom that are not present in the standard expected utility formula and additivity axiom. While this may not seem like such a big difference, it has fairly large effects on the complexity of other probabilistic concepts. For example, conditional believability rankings would be related to unconditional believability rankings in a fairly complex way. Rather than the simple ratio relationship we have with conditional probabilities (where $\text{pr}(A|B) = \text{pr}(A \& B)/\text{pr}(B)$), conditional believability rankings would be fairly complex (specifically, $b(A|B) = 9[(b(A \& B) - 1)/$

26. It is possible to construe someone as a “realist” who takes the “truth” of claims about subjective probabilities as meaningful only within a particular linguistic framework that is adopted on pragmatic grounds (with choices between which is the “true” framework being meaningless). (Compare this to Carnap’s (1950) distinction between “internal” and “external” ontological questions; more on this point below.)

($b(B) - 1$] + 1). Whether it would be the case that *any* alternative definition of degree of belief would be more complex than probabilities, and whether simplicity is indeed relevant to truth (as opposed to usefulness), I will leave open. The main point is clear enough, namely, that for the realist the argument for probabilism from the representation theorems is not complete but has to be supplemented by further arguments, such as arguments of the sort just outlined showing the superiority of the standard approach over other approaches that can also mirror the preference orderings that conform to the axioms of expected utility theory.

Let us suppose that rather than arguing this way, the weak realist decides to take strategy (2), and argues that Maurice's definition of degree of belief is not "significantly" different from the standard one in those respects that are important to attributing degree of belief. One might point out that Maurice's believability rankings are simply linear transformations of Leonard's subjective probabilities, and argue that in the case of probabilities (like utilities and temperatures) this is a difference that makes no difference. This approach commits the weak realist to taking as real properties of Maurice's degrees of belief at most those properties that are common to *both* definitions of degree of belief. (This approach resembles the one discussed earlier in which vague probabilities are defined in terms of the properties that are common to all precise probability functions in a person's representor, except that we are now applying the idea not just to numerical values, but also formal axiomatic properties.) Since Maurice's believability rankings are subadditive rather than additive, this would commit the weak realist to the view that additivity (as defined earlier) cannot be taken *literally* as a property common to all rational degrees of belief (although *having an additive representation* could be). Instead, some more general property would have to do. To pursue this strategy consistently, one would have to investigate the properties of all possible definitions of degree of belief that can, when combined in some way with some representation of value, produce a function that is some order-preserving transformation of EU, since Maurice's believability rankings are only *one example* of a alternative *quantitative* definition of degree of belief that can be formulated consistently with the axioms of expected utility theory. Now, there are *qualitative* properties that subjective probabilities and *any* such "believability rankings" would have to have in common.²⁷ We know that if a person's preferences conform to the axioms of expected utility theory, his degrees of belief can be represented as subjective probabilities. We also know from the work of researchers such as Kraft, Pratt, and Seidenberg (1959) and Chateauneuf and Jaffray (1984) that there are nec-

27. Surveys of qualitative theories of degree of belief can be found in Fine 1973 and Fishburn 1986.

essary and sufficient qualitative conditions for a probability ranking having a quantitative subjective probability representation. For example, let $>^*$ mean “is more probable than” and \geq^* “is no less probable than.” Then having an additive representation implies the following conditions, among others.²⁸

Nontriviality. $S >^* \emptyset$, where S is the necessary event.

Nonnegativity. $A >^* \emptyset$, for all events A.

Asymmetry. If $A >^* B$, then it is not the case that $B >^* A$.

Transitivity. If $A >^* B$ and $B >^* C$, then $A >^* C$.

Monotonicity. If A logically implies B, $B >^* A$.

Qualitative Additivity. If $A >^* B$ and $(A \cup B) \cap C = \emptyset$, then $(A \cup C) >^* (B \cup C)$.

Thus, the weak realist could propose that qualitative principles such as these describe the literally true properties of rational degrees of belief.²⁹ (Note that these all hold of Maurice’s believability rankings, including qualitative additivity.) 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.³⁰ Some probabilists will perhaps regard this as a particularly tenuous form of realism, one that is perhaps not very far from antirealism. Indeed, the view is antirealistic about quantitative subjective probabilities per se (degrees of belief are not literally nonnegative, normalized, or additive, in the sense defined earlier, though we can represent them that way, among others), while remaining realistic about degrees of belief in a more abstract, qualitative sense. The concept of degree of belief on this strategy becomes a purely ordinal notion (although it remains the case that rational degrees of belief would always have a cardinal representation).

Thus, I would argue that though the argument from the representation theorems, and the Reality Condition in particular, is incomplete and ques-

28. Chateauneuf and Jaffray assume a strong Archimedean condition that implies the simpler conditions listed below. Their Archimedean condition can be stated as follows. Let I_X be the indicator function of X, and S the necessary event. Then a qualitative probability on algebra A is *Archimedean* iff for every pair of events A, B with $A >^* B$ there exists $n(A,B) \geq 1$ such that (1) $kI_S - n(I_A - I_B) = \sum_i (I_{C_i} - I_{D_i})$ with $k \geq 0$, $n \geq 1$, and i finite implies (2) $k/n > 1/n(A,B)$. Chateauneuf and Jaffray show that where \geq^* is a complete binary relation on a countable algebra, there exists a probability measure $p(\cdot)$ agreeing with $>^*$ iff \geq^* is Archimedean.

29. Of course, a full implementation of this strategy would still require a defense of the Rationality Condition as well, a task that is beyond the scope of this paper.

30. So, on this view rational degrees of belief are not to be thought of as literally “additive,” in the quantitative sense, though they have an additive quantitative model (subjective probabilities) that can be justifiably accepted on pragmatic grounds.

tion-begging as originally stated, it could be strengthened by either supplementation (by showing that linking degrees of belief to preferences in the standard way is superior to other ways of doing so) or reinterpretation (by showing that there are properties that degrees of belief would have to have however one chose to link degree of belief, utility, and preference), without yielding to either antirealism or eliminativism about degrees of belief.

In conclusion, I would like to address the issue of which of the probabilist views outlined above (antirealism, one of the two versions of weak realism, or a version of strong realism) I regard as preferable. To begin with, I regard the theory of subjective probability (and decision theory as a whole) to be a precise and rigorous development of our folk psychological notions of (degree of) belief, (degree of) desire, confidence, and so on. Although there are eliminativists with respect to folk psychology in general (as well as for subjective probability and utility in particular), I would agree with those who argue that our everyday folk psychological concepts are indispensable to our concepts of personhood and agency (which doesn't of course forbid us from developing them more precisely and rigorously, as in decision theory).³¹ For this reason, I regard the option of antirealism about states such as belief and desire (including their rigorous development in decision theory) as unattractive; though it could arguably be reasonable for a person aware of the scientific evidence for atoms to be an antirealist about them (e.g., if the person wants to limit his or her ontological commitments as much as possible by believing only in the empirical adequacy of theories about atoms), it seems unreasonable for a person to be an antirealist *about those very features that contribute essentially toward making him or her a person and agent* (such as having beliefs, desires, and other contentful mental states). Thus, I would argue that only a form of realism about such mental states will do. Now, with respect to probabilism, and the developed notions of subjective probability and utility, it seems to me that strong realism of the sort defined earlier is more than a probabilist should commit himself to. I do not see any compelling reason for probabilists to think of degrees of belief or utilities as illata appearing as elements in a cognitive or computational system, in the manner of Figure 2. In fact, I would argue that quite the opposite is most likely true. The much-discussed holism of belief-desire attribution (which applies in the context of decision theory as in everyday folk psychology) lends credence to the view that degrees of belief and desire are high-level abstract properties that must be attributed to cognitive systems as a whole. If de-

31. The literature on the adequacy of folk psychology is huge, and I cannot begin to discuss the matter here. See Horgan and Woodward 1985 for a good defense of the indispensability of folk psychological concepts.

degrees of belief, desire, and preference were illata (as described by strong realism), one should in principle be able to identify or conceive of them in isolation (as one can conceive of an atom which makes up a body as existing in isolation;³² one cannot conceive of a body's center of gravity existing apart from the body). Consequently, the relationship between degree of belief, desire, and preference would be wholly contingent; psychological investigation, not definition, would be the appropriate way of determining how they are related. However, I would submit that the relationship is not contingent in this way: if someone claimed that they believed p more strongly than not, preferred A to B , and also preferred (B if p , A if not) to (A if p , B if not),³³ then their claim (or the attribution of one of their preferences) would have to be *wrong*. For reasons such as these, I would agree with Maher (and others) that "attributions of probability and utility [are] essentially a device for interpreting a person's preferences" (Maher 1993, 9).

What, then, about weak realism on strategy (1)? In this case, I would argue that the "theoretical virtues" (simplicity, consilience, fruitfulness, and so on) can only be epistemically relevant to one theory being *more likely* than another when those two theories are not *necessarily* empirically equivalent.³⁴ When two theories are *necessarily* empirically equivalent, then adopting one or the other can only be justified on pragmatic grounds. However, Maurice's and Leonard's self-descriptions necessarily account for the same preferences equally well.³⁵ Consequently, I would opt for the version of weak realism along the lines of strategy (2). People really have degrees of belief, and so belief can be more or less intense, and should (literally) have the qualitative properties stated earlier (such as being asymmetric, transitive, monotonic, and having the property of qualitative additivity in the sense defined above). Moreover, people should have degrees of belief of a sort that can be represented by subjective probability func-

32. This is possible only if the atom is not in an "entangled" quantum state with other atoms in the body.

33. This assumes that the person has no preference whether p obtains or not.

34. This is of course a claim that would take an entire paper to defend. In brief, I would regard as worth consideration at least a justification of simplicity as epistemically relevant to the relative likelihood of theories on the grounds that simpler theories have tended to be more successful empirically than complex ones; but this justification would not hold if the two theories were necessarily empirically equivalent.

35. Note that this line of reasoning only applies in the context of weak realism. If degrees of belief and desire and preference were psychologically distinct, as strong realism says, and each had a distinct physical implementation, then it might be possible (in principle, at least) to determine empirically how they are truly related (just as one could take apart a modular design as represented in Figure 2 and figure out what each of its parts is doing).

tions. This is the realistic part of my view. I would hold further that we are justified in *accepting* (in a sense analogous to van Fraassen's) expected utility theory, of which quantitative subjective probability theory is a part, as a model of opinion, value, and decision, on the grounds that it is well established, elegant, simple, etc. This is the antirealistic part of my position. Subjective probability theory on this view constitutes a useful and compelling model of rational degrees of belief, but not every feature of subjective probabilities can justifiably be understood in the strongly realistic manner of Figure 2. Subjective probabilities "exist" (in the sense that attributions of subjective probabilities to people are acceptable³⁶ if their preferences conform to the axioms of expected utility theory), but only relative to a decision to adopt the usual expected utility framework over other logically possible quantitative frameworks that can also represent preferences of the sort consistent with the axioms of expected utility theory (such as Maurice's alternative proposal). This decision is not wholly justified by the representation theorems alone, but in part on pragmatic or conventional grounds.³⁷

REFERENCES

- Allais, Maurice (1952), "The Foundations of a Positive Theory of Choice Involving Risk and a Criticism of the Postulate and Axioms of the American School", in Maurice Allais and Ole Hagen (eds.) (1979), *Expected Utility Hypotheses and Allais' Paradox*. Dordrecht: Reidel, 27–145.
- Broome, John (1991), *Weighing Goods*. Cambridge, MA: Blackwell.
- Carnap, Rudolf (1950), "Empiricism, Semantics, and Ontology", *Revue Internationale de Philosophie* 4: 20–40. Reprinted in Rudolph Carnap, *Meaning and Necessity: A Study in Semantics and Modal Logic*, 2nd ed., enlarged. Chicago: University of Chicago Press, 1956.
- Chateaufeuf, Alain and Jean-Yves Jaffray (1984), "Archimedean Qualitative Probabilities", *Journal of Mathematical Psychology* 28: 191–204.
- de Finetti, Bruno (1937), "Foresight: Its Logical Laws, Its Subjective Sources", in Henry Kyburg and Howard Smokler (eds.) (1980), *Studies in Subjective Probability*. Huntington, NY: Krieger.
- Dennett, Daniel (1987), *The Intentional Stance*. Cambridge, MA: MIT Press.

36. Again, I mean "acceptable" in the sense of van Fraassen 1980.

37. Several people have commented to me that the view expressed here bears some similarity to the view about "internal" and "external" ontological questions expressed in Carnap 1950. This is to some extent true. Given my commitment to weak realism, I do not regard the question about whether Maurice is "really" an expected utility maximizer with subjective probabilities or a valuation maximizer with believabilities as anything more than a pragmatic question about which language to adopt in describing him. However, there are differences, too: I am not endorsing Carnap's global (and generally anti-metaphysical) application of this distinction. Moreover, I would regard the qualitative features common to the quantitative measures put forward by Leonard and Maurice (and other equally good measures) as real; this is not precisely analogous to the relationship between different linguistic frameworks as Carnap presents it. It also distinguishes me from the antirealist.

- . (1998), *Brainchildren: Essays on Designing Minds*. Cambridge, MA: MIT Press.
- Earman, John (1992), *Bayes or Bust?* Cambridge, MA: MIT Press.
- Fine, Terrence (1973), *Theories of Probability*. New York: Academic Press.
- Fishburn, Peter (1981), “Subjective Expected Utility: A Review of Normative Theories”, *Theory and Decision* 13: 139-199.
- . (1982), *The Foundations of Expected Utility*. Dordrecht: Reidel.
- . (1986), “The Axioms of Subjective Probability”, *Statistical Science* 1: 345-355.
- . (1988), *Nonlinear Preference and Utility Theory*. Baltimore: John Hopkins.
- Gärdenfors, Peter and Nils-Eric Sahlin (eds.) (1988), *Decision, Probability, and Utility*. Cambridge: Cambridge University Press.
- Harman, Gilbert (1986), *Change in View*. Cambridge, MA: MIT Press.
- Horgan, Terrence and James Woodward (1985), “Folk Psychology Is Here to Stay”, *Philosophical Quarterly* 94, no. 2.
- Howson, Colin and Peter Urbach (1989) *Scientific Reasoning: The Bayesian Approach*. La Salle, IL: Open Court.
- Hurley, S.L. (1989), *Natural Reasons: Personality and Polity*. Oxford: Oxford University Press.
- Jeffrey, Richard (1983), *The Logic of Decision*, 2nd ed. Chicago: University of Chicago Press. [First edition, 1965.]
- Jeffrey, Richard (1988), “Conditioning, Kinematics and Exchangeability”, in Brian Skyrms and William Harper (eds.), *Causation, Chance, and Credence*, vol. 1. Boston: Kluwer Academic Press, 221-255.
- Kaplan, Mark (1996), *Decision Theory as Philosophy*. Cambridge: Cambridge University Press.
- Kraft, Charles, John Pratt, and A. Seidenberg (1959), “Intuitive Probability on Finite Sets”, *Annals of Mathematical Statistics* 30: 408-419.
- Maher, Patrick (1993), *Betting on Theories*. Cambridge: Cambridge University Press.
- Pearl, Judea (1988), *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. San Mateo, CA: Morgan Kaufmann Publishers.
- Pearl, Judea and Glenn Shafer (eds.) (1990), *Readings in Uncertain Reasoning*. San Mateo, CA: Morgan Kaufmann Publishers.
- Ramsey, Frank (1926), “Truth and Probability”, in Frank Ramsey, *Philosophical Papers*, edited by D. H. Mellor. Cambridge: Cambridge University Press, 1990, 52-109.
- Savage, Leonard (1954), *The Foundations of Statistics*. New York: Wiley and Sons. [Revised and enlarged edition, New York: Dover, 1973.]
- Skyrms, Brian (1984), *Pragmatics and Empiricism*. New Haven: Yale University Press.
- . (1987), “Coherence”, in Nicholas Rescher (ed.), *Scientific Inquiry in Philosophical Perspective*. Pittsburgh: University of Pittsburgh Press, 225-242.
- van Fraassen, Bas (1980), *The Scientific Image*. New York: Oxford University Press.
- . (1984), “Belief and the Will”, *Journal of Philosophy* 81: 235-256.
- . (1989), *Laws and Symmetry*. New York: Oxford University Press.
- Zynda, Lyle (1996), “Coherence as an Ideal of Rationality”, *Synthese* 109: 175-216.