Radical Probabilism Revisited

Lyle Zynda†‡

In this essay, I analyze and critique Richard Jeffrey's radical probabilism. The basic theses defining it are examined, particularly the idea that probabilistic coherence involves a kind of "consistency." The main challenges to Jeffrey's view are (1) that there is an inconsistency between regarding probabilities as subjective and some probabilistic judgments as better than others, and (2) that decision theory so conceived has no normative import. I argue that both of these challenges can be met.

1. Introduction. Richard Jeffrey developed his "radical probabilism" steadily and consistently over his career, starting from his argument that scientists should not accept hypotheses ("Valuation and Acceptance of Scientific Hypotheses," 1956) and a defense of the concept of probable knowledge ("Probable Knowledge," 1968).[1] During these early years, he also developed his highly significant new approach to decision theory, the "logic of decision," which builds on the subjectivist approach of Ramsey and de Finetti.[2] The central theses of Jeffrey's radical probabilism can be summarized as follows.[3]

1. All-or-nothing belief and desire cannot adequately characterize people's opinions and values; belief and desire both come in degrees.

†To contact the author, please write to: Indiana University South Bend, Philosophy Department, 1700 Mishawaka Ave., South Bend, IN 46634; e-mail: lzynda@iusb.edu.

‡I thank my fellow participants on the panel, Brian Skyrms and Persi Diaconis, and especially Alan Hájek, as well as the audience. It was a splendid session and an honor to participate. Also, thanks to Brad Armendt for helpful lunchtime conversation before the presentation.

1. Both are reprinted in Jeffrey (1992).

2. Jeffrey (1983).

3. Essays 1–7 in Jeffrey (1992) are most relevant to the characterization that follows, particularly "Introduction: Radical Probabilism" (Essay 1), "Probability and the Art of Judgment" (Essay 4, originally published 1985), and "Conditioning, Kinematics, and Exchangability" (Essay 7, 1988). I also rely on Jeffrey (2004) as a final statement of his views and (of course) Jeffrey (1983).

Philosophy of Science, 73 (December 2006): 969–980. 0031-8248/2006/7305-0043$10.00

2. People's degrees of belief and desire are determined by their preferences; the precise way in which this is so is specified in full generality by representation theorems.
3. Degrees of belief, when consistent, conform to the laws of probability theory.
4. All probabilities are judgmental (aka subjective, personal) probabilities; the proper way to characterize "objective" chance and the difference between correlation and causation is in terms of differences in how personal probabilities behave.[4]
5. There is no need to ground epistemology on any certainties, including certainties supposedly given by experience; opinion can be probabilities "all the way down."
6. Observation does not typically deliver a proposition that summarizes all that is learned as a result of it; more typically, observation affects how we redistribute our personal probabilities across a partition, our expectations of the value of random variables, or our odds.
7. When conditional probabilities on members of some partition do not change, the only consistent method of updating one's opinion is probability kinematics (aka "Jeffrey conditioning"); simple conditioning is a special case that occurs when one member of a partition receives probability 1.
8. Other methods of updating one's probabilities can be appropriate when conditional probabilities do not remain stable on any partition.
9. Even though information gained is not always propositional, it can be communicated and summarized by the way in which the observation affected one's odds; Bayes factors are a useful device for doing this.
10. Observation affects our personal probabilities as the result of mostly unconscious processes that ground our judgments, which can be more or less skilled; there is an "art of judgment." Not all personal judgments are equally good, even if consistent.

2. Consistency and Sufficiency in Radical Probabilism. One of the basic themes in Jeffrey's thought is that the only formal constraints on opinion are those of consistency. When consistent, degrees of belief will obey the laws of probability. What is meant by "consistency" is that a person whose beliefs do not conform to the laws of probability will evaluate two logically

4. This is a feature of Jeffrey's solution to the Newcomb problem; he thus maintains his "evidential" decision theory against its "causal" challengers. See Jeffrey (2004, 103–113).

equivalent descriptions of the same situation differently. Thus, it makes sense to speak of "probability logic."

The famous Dutch book argument for the probability calculus Jeffrey takes as dramatizing an underlying inconsistency.[5] Take, for example, the Dutch book argument for the additivity axiom. Suppose that A and B are logically exclusive. Then a bet that pays \$1 if either A or B is true, and nothing otherwise, is equivalent to a book containing two bets, one of which pays \$1 if A obtains, nothing otherwise, and the other of which pays \$1 if B obtains, nothing otherwise. If you regard the payoffs as the only things relevant to their value for you, then the following is a logical inconsistency:

(1) Ticket 1 pays \$1 if either A or B is true, nothing otherwise.
(2) Ticket 2 pays \$1 if A is true, nothing otherwise.
(3) Ticket 3 pays \$1 if B is true, nothing otherwise.
(4) The possible payoffs of tickets 1–3 are a measure of their value, and only this is relevant to their value.
(5) Ticket 1 is more (or less) valuable than tickets 2 & 3 together.

If A and B are logically exclusive, (1)–(5) are logically inconsistent; since (1)–(4) logically imply that ticket 1 has the same value as tickets 2 & 3 together: thus, whatever you'd pay for tickets 2 & 3 together, this has to equal the price you'd pay for ticket 1 alone. Identifying your probability with the price you regard as fair for each ticket, your probabilities must be additive.

(6) Your probability for X is p iff you regard a price of \$$p$ as fair for a bet that pays \$1 if X, nothing otherwise.

Note that Jeffrey does not hold that such price setting during betting is an operational definition of one's personal probabilities. Jeffrey agrees that in general the price that one will pay for tickets need not reflect one's probabilities, that there are many propositions on which one cannot bet, and so on. Betting is a very special case of decision making under conditions of risk. As is well known, there are many objections to Dutch book arguments in the literature, which focus on limitations of the Dutch book scenario. Jeffrey's strategy to counter these is as follows. He notes that it is unreasonable to suppose that probability follows different rules for different circumstances. Why, he asks, should one's probabilities obey one set of rules in situations where the value of the bets for you *is* measured well by their monetary outcomes (a situation he supposes to obtain at least sometimes), and a different set of rules in other situations? There

5. See, e.g., Jeffrey (2004, 4–9).

should be one set of rules that governs degrees of belief in all contexts. Therefore, if you can show (as the Dutch book argument does) that the axioms of probability theory must hold if one is to be consistent in a certain *type* of situation (the Dutch book situation), then the laws of probability should hold in *all* situations. Always following the same rules is, of course, a kind of "consistency"—though pragmatic, rather than logical.

Ultimately, however, since (6) cannot plausibly be regarded as a definition, the definition of personal probabilities rests on a rationale that identifies them with the mathematical expectation of utility. Dutch book arguments are simply useful dramatizations. One's degrees of belief (and desire) are defined by one's preferences. The precise way in which this is so is given in full generality by representation theorems. Thus, ultimately the justification for probabilism rests on the representation theorems.[6] These say that if preferences meet certain conditions, there is a representation of those preferences in terms of personal probabilities and desirabilities—that is, a function of those probabilities and desirabilities that ranks items exactly as the preferences do. The conditions on preferences vary from theory to theory, but it is typically supposed that some of these are *rationality* conditions. (Others might be there for technical convenience, e.g., to define a real-valued function to represent strength of preference.)

Jeffrey's emphasis on consistency extends to updating rules. Similar to the synchronic case, diachronic Dutch book arguments (such as David Lewis's famous one) dramatize an underlying inconsistency. The idea is that if you judge that certain observations will have certain effects on your view as you update, this implies (mathematically) that certain types of updating rules apply. For example, it is mathematically equivalent to say that (*a*) you update by Jeffrey conditioning from probability functions P to Q and (*b*) a partition exists on which your conditional probabilities (as defined by P and Q) do not change. Therefore, if you say that your conditional probabilities will remain the same on a certain partition, and also that you will not update by Jeffrey conditioning, you are being mathematically inconsistent.

> **Sufficiency Principle**: A change from probability functions P to Q occurs by Jeffrey conditioning (defined as $Q(p) = \sum_i (Q(e_i)P(p|e_i)$ for all p and some partition $\{e_i\}$) iff for all p and e_i, $P(p|e_i) = Q(p|e_i)$).

6. Jeffrey's longtime friend and collaborator Brian Skyrms also defends this view. See, e.g., Skyrms (1987).

(In what follows, I will refer to this equivalence as the Sufficiency Principle, and the clause following the biconditional, stated informally as (*b*) above, as the *sufficiency condition* or just *sufficiency*.[7])

According to Jeffrey (1992, 122), diachronic Dutch book arguments (such as Lewis's) implicitly assume that conditional probabilities remain fixed: "Where does the sufficiency condition enter Lewis' argument? It enters with the assumption that today, while your probability distribution is P, you know what tomorrow's Q will be if you learn that D is true. Then your odds between ways in which D might come true are determined by today's judgments: by P. That's why tomorrow's odds between propositions that imply D are the same as today's. That's where the sufficiency condition enters."

What Jeffrey effectively presents us with is a trilemma, which holds as a matter of mathematical necessity.

> **Trilemma**. Either (1) the change from P to Q occurs by Jeffrey conditioning on $\{e_i\}$; or (2) probabilities conditional on members of $\{e_i\}$ do not remain fixed from P to Q for at least one proposition; or (3) either P or Q is not a probability function.

This trilemma is, I would submit, the basis of Jeffrey's argument for probability kinematics. It too rests on a kind of consistency. It is important to note here that Jeffrey is not claiming that you always ought to update by Jeffrey conditioning (of which simple conditioning is a special case). It is only appropriate to Jeffrey condition if your conditional probabilities remain fixed over some partition. In fact, if your conditional probabilities do not remain so fixed, it is *inconsistent* to Jeffrey condition; indeed, it is *mathematically impossible*. This raises a puzzle: we have seen that Jeffrey conditioning is not always applicable (if sufficiency fails), but when sufficiency holds, you seemingly have *no choice* but to do it (because of the mathematical equivalence stated in the Sufficiency Principle). So, where's the inconsistency, if you can't actually do anything that violates the trilemma? In fact, there are only two possibilities: (*a*) there is a synchronic inconsistency either before or after the change (item 3 of the trilemma is violated), or (*b*) there are inconsistencies in one's *intentions* (one intends to do the impossible, i.e., to hold fixed probabilities on some partition without Jeffrey conditioning).

When should the sufficiency condition hold? Ultimately, Jeffrey argues, whether one's conditional probabilities should remain fixed on some partition is a matter of personal judgment, just like one's synchronic proba-

7. The equivalence stated in the Sufficiency Principle only holds when P and Q are both probability functions. Thus, the Principle assumes synchronic coherence.

bilistic judgments. To examine Jeffrey's views on the nature of such judgments, let us begin by looking at a simple example (Jeffrey 2004, 3–4).

> About 49% of observed births are girls, 51% boys. Given this, what's the probability the first baby born tomorrow in this city will be a boy?
>
> Let a goy be a girl born before midnight tonight, or a boy born afterward. About 49% of observed births have been goys, the other 51% birls. Given this, what's the probability the first baby born tomorrow in this city will be a birl?

One response to this might be that the answer is objectively determined by which of the predicate pairs "boy/girl" or "goy/birl" is *projectible*. Jeffrey resists this. He says merely that which predicate pair is projectible is a matter of judgment, that is, you can only consistently regard one pair as projectible, and which one you'd project determines which is projectible *for you*. (In fact, all except the perverse will project the boy/girl pair.)

Similarly, judgments of equiprobability are also just that—probabilistic judgments. One cannot "divide cases" into "equally probable" parts without making some kind of personal judgment.

A similar analysis would apply to other well-known examples, such as the Three Prisoners problem and the structurally similar Monty Hall problem. Let's consider the latter.[8] You are faced with three doors; a prize is behind exactly one of them. After you pick one door (say, door 1), the game host opens one of the remaining two doors, one he knows has no prize behind it, say, door 2. You are now given an opportunity to stick with your original choice (door 1) or switch to the other door remaining closed (door 3). What should you do? (See Figure 1.)

One claim is that you ought to switch. The rationale is that your probability of choosing the correct door was originally 1/3, and that this is the same even after a door without a prize is opened. So, the probability that the remaining closed door has the prize is now 2/3.

If you conceive of yourself as conditioning on the proposition "the prize is behind either door 1 or door 3," however, this is impossible. In that case, you must give probability 1/2 to the prize being behind door 1 once door 2 is opened. You gave even odds between doors 1 and 3 before, so if you condition on the prize not being behind door 2, the odds between them must remain even. Thus, you should be indifferent between sticking with door 1 and switching to door 3.

One explanation for the discrepancy between these two informal modes of reasoning is that each of them obtains under different assumptions.

8. Jeffrey discusses the former in Jeffrey (1992, 122–124).

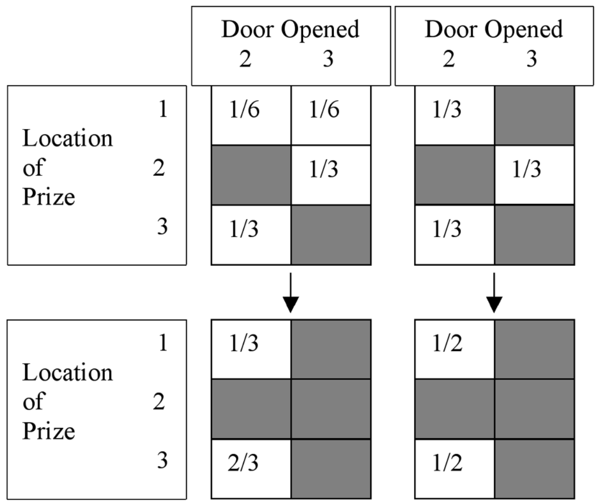

Figure 1.

You need to consider a finer partition, which not only takes into account where the prize is, but how the host chooses which door to open. One thing you should definitely *not* do is simply condition on the proposition that the prize isn't behind door 2, since there is a finer (more informative) partition that is relevant.

The difference between the two distributions on this partition is whether the host will choose which empty door to reveal randomly, when he can (left), or always a particular door (if he has a choice), say, door 2 (right). In the former case, the probability that the prize is behind door 3 once 2 is opened will be 2/3, and I should switch. In the latter, the probabilities that the prize is behind door 1 or 3 will be equal. In both cases, you update by simple conditioning on a finer partition.

Does there always have to be a finer partition to decide such cases? Is Jeffrey conditioning ultimately grounded on simple conditioning? We do know (thanks to van Fraassen) that any case of Jeffrey conditioning can be reproduced by conditioning on a deeper partition, in particular, one about your future probabilities (assuming that you always obey the Reflection principle—a controversial assumption).[9] Also, in some situations (when the ratios of one's new to old probabilities have an upper bound), simple conditioning in a larger algebra (superconditioning) can reproduce

9. See van Fraassen (1980, 1984). Jeffrey discusses Brian Skyrms's version of the proof in Jeffrey (1992, 126–127).

a case of Jeffrey conditioning (Diaconis and Zabell 1982). However, Jeffrey denies that there need always be a deeper partition on which one *really* conditions. You can always *choose* a finer partition, but whether you regard it as *appropriate* is a matter of judgment.

Here is an argument that might be used to bolster Jeffrey's case. Consider the following principle:

> **Internalism**: A certain model of your change in mind is only a justification for that change if the elements of that model were used by you to justify that change.

In other words, if we are thinking about my judgments as I am consciously aware of them, whether they are justified is determined wholly by considerations internal to those judgments. If we accept internalism, it follows that we cannot justify a change from P to Q by conditioning on future opinion or superconditioning unless that is how I conceived of my change and justified it for myself. Another way of looking at this is in terms of cognitive processes; for me to make up my mind about something, there is a certain mental process that leads up to it. That mental process—my conscious deliberation—is determined by the contents of consciousness. Decision theory must relate to what I use to make my decision; it should not merely link the end points of my decision making (where I started and ended). Thus, if I didn't think of my future opinion at all while changing my view from P to Q, then you cannot justify my change in view by saying I *would* have arrived at the same probability distribution Q *had* I conditioned on various possible future opinions (and similarly for a larger partition that I never actually use to supercondition).

If this is the case, then if I judge that conditional probabilities remain fixed on some partition, the elements of which I reassign probabilities, then the fact of the matter is that I have Jeffrey conditioned on that partition; I have *not* conditioned on some deeper partition.

Jeffrey argues that this is actually how we often change our minds.[10] When we have expertise, for example (like a skilled pathologist), what we observe directly affects how we distribute our credence across some partition—and also how probabilities conditioned on members on that partition will change. This is all part of the "art of judgment." To learn how to make such judgments reliably is a skill that must be learned (that is what distinguishes novices from experts). Moreover, cognitive science and neuroscience are increasingly telling us that most of skilled practitioners' judgments (in any trade) result from much *unconscious* processing. (Indeed, some have gone so far as to say that our unconscious brain *always*

10. See Jeffrey (1992, Essay 1).

makes decisions, after which our conscious minds merely become aware of a fait accompli—see, e.g., Wegner 2002.) One purpose of the logic of decision is to provide a means of regularizing our conscious deliberations and to communicate the contents of our skilled judgments to others (e.g., via Bayes factors).

Let's go back to Monty Hall. Suppose that I adopt the switching strategy. Why should I not view this as Jeffrey conditioning on the original partition (the three possible locations of the prize), where my judgment is simply that probabilities conditioned on the partition (door 2, not door 2) change after I learn whether door 2 has the prize (by seeing it opened)? Do I need to justify this "probable knowledge" by appealing to some deeper assumption I allegedly make about how Monty Hall will choose between doors if he has a choice? Perhaps in *this* case I will in fact find this assumption relevant; but if so, that is a feature of my judgments. In other cases, there need be no such deeper assumptions (on finer partitions) relevant to my *actual* deliberations.[11] One simply cannot get rid of judgmental probabilities.

It is worth noting that Jeffrey also resists moving to a deeper partition in his solution to the problem of old evidence (Jeffrey 1992, 103–107, and 2004, 44–47). Rather than conditioning on a proposition in a "finer" partition, such as that h entails e, he imagines an ur-prior in which a contradiction ($h \,\&\, \neg e$) is assigned positive probability, and specifies changes in terms of Bayes factors that lead to an increase in probability for h when it is discovered that h entails e.[12]

3. Possible Tensions in Jeffrey's Radical Probabilism (with Suggested Resolutions). Jeffrey defends a form of subjectivism, which holds that personal probabilities are all that exist and that the only formal constraints on these are consistency. Yet he holds that some probabilistic judgments are better than others. For example, a reasonable person will rely on medical experts rather than faith healers. Why? Well, the obvious answer is that the medical experts get things right more reliably. In other words, there is a better match between their probabilistic judgments and the actual observed frequencies. But what are these "observed frequencies"? If you are a thoroughgoing subjectivist, all you can appeal to is *beliefs* about frequencies. This raises a problem: one consistent way to avoid adjusting one's beliefs is simply to deny the frequencies another person claims to have obtained

11. For example, it may have never occurred to me even to think about how Monty Hall chooses which door to open; thus, I have no opinion about that at all.

12. There is a tension between this solution to the problem of old evidence and Jeffrey's logic of decision, as formally developed, since there an agent's preferences, probabilities, and desirabilities are defined on an atomless Boolean algebra *from which the impossible proposition has been removed.*

(and in fact did obtain). Call this strategy *seeing is disbelieving*. (In realistic cases, this could result from what cognitive psychologists call *confirmation bias*—counting only observations that confirm one's hypothesis, ignoring those that do not.) This is *consistent*; why is it *unreasonable*?

One answer to this might appeal to people's actual dispositions. Given how we are constituted (as normal human beings), we cannot (if sane) deny that a coin came up heads when it appears to do so in bright light right in front of us. However, even if this is so, *interpretation* is required in many other cases to identify something as an "occurrence" of some event type. (Is this an instance of seeing a mobile chemical weapons unit, or just a hydrogen refilling station? Is this an instance of a failure of President Bush's foreign policy, or not?) The question thus arises whether there are reasonable and unreasonable interpretations—each of which may be self-consistent.

Moreover, consistency does not require that one's conditional probabilities be one way rather than another (barring cases of logical implication). All we can say is: if one's conditional probabilities are such and such, and if they do not change upon observations, then consistency requires updating by Jeffrey conditioning. Suppose that someone chooses to maintain a belief in a faith healer's powers simply by *explaining away* observed failures, for example, by saying they don't count against him. In probabilistic terms, this means that this person's probabilities for the faith healer having powers are independent of the observed frequencies to date.

$$
\begin{aligned}
&P(\text{the faith healer has powers to heal} \\
&\quad |x\% \text{ of previous attempts at healing were successful}) \\
&= P(\text{the faith healer has powers to heal}).
\end{aligned}
$$

Call this the *dogmatic* strategy. (The equation would hold, for example, if P (the faith healer has powers to heal) $= 1$.) One might also call this *believing is seeing*. We know in fact that people often take a dogmatic stance toward certain things and become impervious to any new evidence. Since a dogmatist can be consistent, this means that we cannot always criticize a dogmatist based on considerations of consistency alone.

Where does this leave us? I would submit that there is in fact no contradiction between Jeffrey (or those of us who agree with him) saying that all probability judgments are personal ("subjective") and that some judgments are better than others. These are consistent if one holds that there are different *levels* of criticism: (1) consistency conditions and (2) informal conditions that go beyond these (such as "your probabilities should match the actual frequencies insofar as it is possible," "seek disconfirming as well as confirming evidence," "avoid wishful thinking," or even "don't give probability 1 to a contingent proposition").

To hold that one's degrees of belief and desire are determined by one's preferences creates a second tension that some might argue undermines the status of the logic of decision as *normative*. If preferences alone mathematically *determine* degrees of belief and desire (as specified by an appropriate representation theorem), degrees of belief and desire cannot exist apart from those preferences, nor can they conflict with those preferences. The picture that results is of preferences and degrees of belief and desire necessarily changing *together*. This is different from the picture where one *uses* one's degrees of belief and desire to *form* one's preferences. In the former, degrees of belief and desire are in effect *epiphenomenal*; in the latter, there's a causal process involving distinct interacting mental states (beliefs, desires, and preferences).

Normativity requires that one can *fail* to meet standards, and this requires the causal interaction picture. For example, suppose that according to some representation theorem, your preferences require that your degree of belief in p is 1/2; yet it is actually different. Or, you might believe p to degree 1/2, and want X more than Y, and believe that if you perform a certain act A, X will obtain if p and Y if not; and yet not have formed a preference between A and the status quo. In this case, a causal process of deliberation will *form* the preference and not simply reveal to you what you *already* prefer. In the former case, you have reason to *revise* the degree of belief (or the preferences that imply it must be 1/2). This requires, however, that in both cases one's degrees of belief *must* be determined by something *other* than one's entire system of preferences. What is this "other" thing? It is implausible to hold that we can always directly introspect our degrees of belief, for (1) we can be wrong about our degrees of belief (self-deception is sometimes possible), and (2) a degree of belief is not a mere disposition to report a number (or qualitative probability judgment, for that matter).

Another example of this difficulty of distinguishing descriptive and normative comes from Jeffrey's view toward updating rules. Apparently, given the Sufficiency Principle is mathematically necessary, you cannot tell someone that he *ought to* update by Jeffrey conditioning, only that he *does* update by Jeffrey conditioning when certain conditions obtain.

My suggested resolution to this second tension is that we (1) admit that degrees of belief, desire, and preferences somehow have to be independently existing, interacting states, but also (2) insist that there is no clear line between the descriptive and the normative. What (2) means is that we regard degrees of belief as sometimes locally determined (perhaps vaguely), by a small number of preferences, without precluding global criticism of these vis-à-vis our preferences as a whole. Thus, both to attribute and to critique degrees of belief involves working toward a kind of reflective equilibrium between the small (local degrees of belief, desire, and preferences) and the large (entire systems of these).

More precisely, the objection that Jeffrey's radical probabilism is not of normative significance can be deflected if we maintain that degrees of belief and desire can be attributed to an agent if their preferences are at least locally coherent. In other words, there must be *some* preferences (an incomplete set) that do not contradict the assumptions of a representation theorem, if any attribution (of degrees of belief or desire) is to be possible.[13] However, this may not hold *globally* of the person's preferences. This is another way for a person to be "inconsistent" (hence, criticizable). Or, their preferences (globally speaking) may not be completely defined (in which case they can form *new* preferences). Alternatively, the person's preferences might not meet the normatively appropriate assumptions behind a representation theorem—and so we do not have a description of their degrees of belief and desire based on *it*—but it might meet the conditions of some more general representation theorem that leaves out one or more rationality conditions. (For example, the person might violate the "sure thing" principle, but still have asymmetric, transitive preferences.[14]) In this sense, a person can be said to "have" degrees of belief based on the more general representation theorem. Obviously, what needs to be investigated are the limits on such generalizations of alternative preference theories giving us adequate *descriptions* of degrees of belief and desire. Meanwhile, Jeffrey's theory can be defended as normative since it allows for *local* criticism of defective degrees of belief, desire, and preferences.

REFERENCES

Diaconis, Persi, and Sandy Zabell (1982), "Updating Subjective Probability," *Journal of the American Statistical Association* 77: 822–830.

Fishburn, Peter (1988), *Nonlinear Preference and Utility Theory*. Baltimore: Johns Hopkins University Press.

Jeffrey, Richard (1983), *The Logic of Decision*. 2nd ed. Chicago: University of Chicago Press. 1st ed., 1965.

——— (1992), *Probability and the Art of Judgment*. Cambridge: Cambridge University Press.

——— (2004), *Subjective Probability: The Real Thing*. Cambridge: Cambridge University Press.

Skyrms, Brian (1987), "Coherence," in N. Rescher (ed.), *Scientific Inquiry in Philosophical Perspective*. Pittsburgh: University of Pittsburgh Press, 225–242.

van Fraassen, Bas (1980), "Rational Belief and Probability Kinematics," *Philosophy of Science* 47: 165–187.

——— (1984), "Belief and the Will," *Journal of Philosophy* 81: 235–256.

Wegner, Daniel (2002), *The Illusion of Conscious Will*. Cambridge, MA: MIT Press.

13. In such a case, we can use the set of all probability-utility pairs defined by some complete extension of the local, incomplete preferences to represent their (locally determined) degrees of belief.

14. See Fishburn (1988) for a survey of nonstandard preference theories.