

A Nonpragmatic Vindication of Probabilism*

James M. Joyce^{†‡}

Department of Philosophy, University of Michigan

The pragmatic character of the Dutch book argument makes it unsuitable as an “epistemic” justification for the fundamental probabilist dogma that rational partial beliefs must conform to the axioms of probability. To secure an appropriately epistemic justification for this conclusion, one must explain what it means for a system of partial beliefs to accurately represent the state of the world, and then show that partial beliefs that violate the laws of probability are invariably less accurate than they could be otherwise. The first task can be accomplished once we realize that the accuracy of systems of partial beliefs can be measured on a *gradational* scale that satisfies a small set of formal constraints, each of which has a sound epistemic motivation. When accuracy is measured in this way it can be shown that any system of degrees of belief that violates the axioms of probability can be replaced by an alternative system that obeys the axioms and yet is more accurate *in every possible world*. Since epistemically rational agents must strive to hold accurate beliefs, this establishes conformity with the axioms of probability as a norm of *epistemic* rationality whatever its prudential merits or defects might be.

1. Introduction. According to the doctrine of *probabilism* (Jeffrey 1992, 44) any adequate epistemology must recognize that opinions come in

*Received November 1997; revised February 1998.

†Send requests for reprints to the author, Department of Philosophy, University of Michigan, 435 South State Street, Ann Arbor, MI 48109–1003.

‡I have been helped and encouraged in the development of these ideas by Brad Armendt, Robert Batterman, Alan Code, David Christensen, Dan Farrell, Allan Gibbard, Alan Hajek, William Harper, Sally Haslanger, Mark Kaplan, Jeff Kasser, Louis Loeb, Gerhard Nuffer, Peter Railton, Gideon Rosen, Larry Sklar, Brian Skyrms, Bas van Fraassen, David Velleman, Peter Vranas, Nick White, Mark Wilson, Steve Yablo, and Lyle Zynda. Richard Jeffrey’s influence on my thinking will be clear to anyone who knows his writings. Special thanks are also due to two anonymous referees from *Philosophy of Science*, whose splendidly detailed comments greatly improved the final version of this paper.

Philosophy of Science, 65 (December 1998) pp. 575–603. 0031-8248/98/6504-0002\$2.00
Copyright 1998 by the Philosophy of Science Association. All rights reserved.

varying gradations of strength and must make conformity to the axioms of probability a fundamental requirement of rationality for these *graded* or *partial* beliefs.¹ While probabilism has long played a central role in statistics, decision theory, and, more recently, the philosophy of science, its impact on the traditional theory of knowledge has been surprisingly modest. Most epistemologists remain committed to a *dogmatist* paradigm that takes *full belief*—the unqualified acceptance of some proposition as true—as the fundamental doxastic attitude. Partial beliefs, when considered at all, are assigned a subsidiary role in contemporary epistemological theories.

Probabilism's supporters deserve part of the blame for this unhappy state of affairs. We probabilists typically explicate the concept of partial belief in pragmatic terms, often quoting Frank Ramsey's dictum that, "the degree of a belief is a causal property of it, which we can express vaguely as the extent to which we are prepared to act on it" (1931, 166). Moreover, when called upon to defend the claim that rational degrees of belief must obey the laws of probability we generally present some version of the *Dutch Book Argument* (Ramsey 1931, de Finetti 1964), which establishes conformity to the laws of probability as a norm of *prudential* rationality by showing that expected utility maximizers whose partial beliefs violate these laws can be induced to behave in ways that are sure to leave them less well off than they could otherwise be. This overemphasis on the pragmatic dimension of partial beliefs tends to obscure the fact that they have properties that can be understood independently of their role in the production of action. Indeed, probabilists have tended to pay little heed to the one aspect of partial beliefs that would be of most interest to epistemologists: namely, their role in representing the world's state. My strong hunch is that this neglect is a large part of what has led so many epistemologists to relegate partial beliefs to a second-class status.

I mean to alter this situation by first giving an account of what it means for a system of partial beliefs to accurately represent the world, and then explaining why having beliefs that obey the laws of probability contributes to the basic epistemic goal of accuracy. This strategy is not new. Roger Rosenkrantz (1981) has taken a similar approach, arguing that if the accuracy of degrees of belief is measured by a quantity called the *Brier score*, then systems of degrees of belief that violate the laws of probability are necessarily less accurate than they need to be. In a similar vein, Bas van Fraassen (1983) and Abner Shimony

1. A further tenet of the view is that *Bayesian conditioning* is the only legitimate method for revising beliefs in light of new evidence. This aspect of probabilism, which remains an active topic of debate in philosophical circles, will not be our concern here.

(1988) have maintained that accuracy can be measured using a quantity called the *calibration index*, and they have argued, in slightly different ways, that any system of degrees of belief that violates the probability axioms can be replaced by a better calibrated system that satisfies them. While both these approaches are on the right track, we shall see below that neither ultimately succeeds. The van Fraassen/Shimony strategy fails because calibration is not a reasonable measure of accuracy for partial beliefs, and Rosenkrantz ends up begging the question (albeit in a subtle and interesting way).

To secure my nonpragmatic vindication of probabilism I will need to clarify the appropriate *criterion of epistemic success* for partial beliefs. The relevant success criterion for full beliefs is well-known and uncontroversial.

The Norm of Truth (NT):² An epistemically rational agent must strive to hold a system of full beliefs that strikes the best attainable overall balance between the epistemic good of fully believing truths and the epistemic evil of fully believing falsehoods (where fully believing a truth is better than having no opinion about it, and having no opinion about a falsehood is better than fully believing it).³

2. Even though the Norm of Truth is widely accepted, there is no consensus about the basis of its prescriptive force. Some read it as expressing a prima facie intellectual obligation that is binding on all believers (Chisholm 1977, 7). Others portray it as an “internal” norm that is partially constitutive of what it is to *be* a believer, so that an attitude toward X cannot even be counted as a full belief (as opposed to a supposition or wish that X) unless its holder is committed to regarding the attitude as successful iff X is true. See, e.g., Anscombe 1957, Smith 1987, and Velleman 1996. A third view, which has been championed by Richard Foley (1987, 66), sees the Norm as being grounded in our practices of epistemic evaluation; terms like “justified” or “epistemically rational” can only be meaningfully applied to individuals who regard their full beliefs as successful iff they are true. For present purposes, it does not matter which of these rationales for the Norm of Truth one adopts. The important point is that there is little real dispute about its status as a basic criterion of epistemic success for full beliefs.

3. Mark Kaplan has observed that the Norm of Truth is not a pure accuracy principle since it places a premium on believing truths as against suspending judgment. He suggests, however, that none of my arguments rely upon this aspect of Norm, and that I could have just as easily made accuracy for systems of full belief a matter or the their truth-to-falsehood *ratio*. While I think this is right, I have decided to stick with NT as my official “success condition” for full beliefs because doing so helps make sense of some important debates in the epistemology of full belief. Notice that NT does not say *how much* better (worse) it is to believe a truth (falsehood) than it is to have no opinion about it, nor does it give any hint about what the best overall balance of truths to falsehoods might be. The way we decide these issues will greatly effect the form of dogmatic epistemology. For example, those who tend to put great emphasis on the avoidance of error may see only a small difference between believing truly and suspending belief whereas the difference between suspending belief and believing falsely

This principle underlies much of dogmatic epistemology. It implies that we should aim to accept truths and reject falsehoods whenever we have a choice in the matter, that we should evaluate our full beliefs, even those we cannot help holding, on the basis of their truth-values, and that we should treat evidence for the truth of some proposition as a *prima facie* reason for believing it. Probabilism's main shortcoming has been its inability to articulate any similarly compelling criterion of epistemic success to serve as the normative focus for an epistemology of partial belief. I shall formulate and defend such a criterion, and prove that holding degrees of belief that obey the laws of probability is an essential prerequisite to its satisfaction. This will establish the requirement of probabilistic consistency for partial beliefs as a norm of *epistemic* rationality, whatever its prudential costs or benefits might be.

My argument will be based on a new way of drawing the distinction between full and partial beliefs. The difference between these two sorts of attitudes, I claim, has to do with the appropriate *standard of accuracy* relative to which they are evaluated. While both "aim at the truth," they do so in quite different ways. Full beliefs answer to a *categorical*, "miss is as good as a mile," standard of accuracy that recognizes only two ways of "fitting the facts": getting them exactly right or having them wrong, where no distinctions are made among different ways of being wrong. This is reflected in the Norm of Truth, which is really nothing more than the prescription to maximize the categorical accuracy of one's full beliefs.

A simple accurate/inaccurate dichotomy does not work for partial beliefs because their accuracy is ultimately a matter of degree. As I shall argue, partial beliefs are appropriately evaluated on a *gradational*, or "closeness counts," scale that assigns true beliefs higher *degrees of accuracy* the more strongly they are held, and false beliefs lower *degrees of accuracy* the more strongly they are held. My position is that a rational partial believer must aim not simply to accept truths and reject falsehoods, but to hold partial beliefs that are gradationally accurate by adjusting the strengths of her opinions in a way that best maximizes her degree of confidence in truths while minimizing her degree of confidence in falsehoods. For the same reasons⁴ that a person should aim to hold full beliefs that are categorically accurate, so too should she aim to hold partial beliefs that are gradationally accurate. We thus are lead to the following analogue of the Norm of Truth:

may loom quite large. Conversely, Popperians who want to encourage "bold conjecturing" will emphasize the "believe the truth" aspect of the Norm of Truth and downplay its prescription to avoid the false.

4. The options here are roughly the same as those listed in fn. 2.

The Norm of Gradational Accuracy (NGA): An epistemically rational agent must evaluate partial beliefs on the basis of their gradational accuracy, and she must strive to hold a system of partial beliefs that, in her best judgment, is likely to have an overall level of gradational accuracy at least as high as that of any alternative system she might adopt.

The system of partial beliefs with the highest attainable level of gradational accuracy will, of course, always be the one in which all truths are believed to the maximum degree and all falsehoods are believed to the minimum degree. This does not, however, imply that an epistemically rational agent must hold partial beliefs of only these two extreme types. Indeed, she should rarely do so. Unlike full believers, partial believers must worry about the epistemic costs associated with different ways of being wrong. Since the worst way of being wrong is to be maximally confident in a falsehood, there is a significant epistemic disincentive associated with the holding of extreme beliefs. Indeed, I shall argue that on any reasonable measure of gradational accuracy the incentive structure will force a rational agent to “hedge her epistemic bets” by adopting degrees of belief that are indeterminate between certainty of truth and certainty of falsehood for most contingent propositions.

The Norm of Gradational Accuracy will be the cornerstone of my nonpragmatic vindication of probabilism. To show that epistemically rational partial beliefs must obey the laws of probability, I will first impose a set of abstract constraints on measures of gradational accuracy, then argue that these constraints are requirements of epistemic rationality, and finally explain why conformity to the laws of probability improves accuracy relative to any measure that satisfies them. It will then follow from NGA that it is irrational, from the purely epistemic perspective, to hold partial beliefs that violate the laws of probability.

There are five sections to come. Section 2 sketches a version of the Dutch book argument and explains why it does not provide an appropriately “epistemic” rationale for conforming one’s degrees of belief to the axioms of probability. Section 3 introduces the notion of gradational accuracy and explains why it is the appropriate standard of evaluation for degrees of belief. Section 4 criticizes rival accounts of accuracy for partial beliefs, and presents a formal theory of gradational accuracy. Section 5 shows that degrees of belief which violate the axioms of probability are less accurate than they otherwise could be relative to *any* reasonable measure of accuracy. Section 6 explains how these results can be applied to more realistic cases in which agents are not assumed to have precise numerical degrees of belief.

2. The Dutch Book Argument and its Shortcomings. To specify a partial belief one must indicate a *proposition* X and the *strength* with which it is held to be true. We will imagine that the propositions about which our subject has beliefs are included in a σ -complete Boolean algebra Ω , i.e., a non-empty set of propositions that is closed under negation and countable disjunction. The strength of the person's belief in X is a matter of how confident she is in its truth. For the moment, we will engage in the useful fiction that our agent's opinions are so definite and precise that their strengths can be measured by a real-valued *credence function* b that assigns every proposition $X \in \Omega$ a unique *degree of belief* $b(X)$. This is absurd, of course; in any realistic case there will be many propositions for which a rational agent need have no definite *degree* of belief. We discuss these imprecise beliefs in the last section of the essay.

According to probabilism, a rational believer's credence function must obey the laws of probability:

Normalization: $b(X \vee \sim X) = 1$.

Non-negativity: $b(X) \geq 0$ for all $X \in \Omega$.

Additivity: If $\{X_1, X_2, X_3, \dots\}$ is a finite, or denumerably infinite, *partition* of the proposition X into pairwise incompatible disjuncts, so that $X = (X_1 \vee X_2 \vee X_3 \vee \dots)$ where X_j and X_k are incompatible for all j and k , then $b(X) = b(X_1) + b(X_2) + b(X_3) \dots$

The principal aim of this essay is to provide a justification of the probabilist's "fundamental dogma" that rational agents must have degrees of belief that obey these three laws.

To understand the justification I am going to give, it will be useful to begin by considering a particularly revealing version of the Dutch book argument due to Bruno de Finetti (1974) and Leonard Savage (1971). Even though this argument ultimately fails to provide an acceptable epistemic rationale for the fundamental dogma it does suggest a fruitful way of approaching the problem. De Finetti and Savage developed an ingenious piece of psychometrics, which I call the *prevision game*, that was designed to reveal the strengths of a person's partial beliefs. To simplify things they assumed they were dealing with a *miser* who desires only money, and whose love of it remains fixed no matter how rich or poor she might become.⁵ This miser is presented with a list

5. In saying that a miser loves only money we imply that (a) all her desires are directed toward propositions that specify her net worth under various contingencies, and (b) that money has *constant marginal utility* for her, so that giving her an extra dollar always increases her happiness by the same amount no matter how large her fortune might be. Proponents of the Dutch book do of course realize that no misers

of propositions $X = (X_1, X_2, \dots, X_n)$ and is offered a dollar to designate a corresponding sequence of real numbers $p = (p_1, p_2, \dots, p_n)$. The catch is that she must repay a portion of her dollar once the truth-values of the X_j have been revealed. The size of her loss is fixed by the game's *scoring rule*, a function $S(p, \omega)$ that assigns a penalty of up to \$1 to each pair consisting of a joint truth-value assignment ω for the propositions in Ω (hereafter a "possible world"), and a sequence of numbers p . For reasons that will be made clear shortly, de Finetti and Savage focused their attention on games scored using *quadratic-loss* rules that have the form $S(p, \omega) = \sum_i \lambda_i [\omega(X_i) - p_i]^2$ where $\lambda_1, \dots, \lambda_n$ are non-negative real numbers that sum to one and $\omega(X_i)$ is the truth-value (either 0 or 1) that X_i has at world ω . An illuminating example is provided by the rule that weights each X_i equally, so that $\lambda_1 = \lambda_2 = \dots = \lambda_n = 1/n$. This is called the *Brier score* in honor of the meteorologist George Brier (1950), who proposed that it be used to measure the accuracy of probabilistic weather forecasts (as in, "the chance of rain is 30%"). Following de Finetti, let us call the numbers that an agent reports in a game scored using a quadratic-loss function her *previsions* for the various X_i .

De Finetti and Savage used quadratic-loss functions to score prevision games because these rules have two properties that make them uniquely suited to the task. First, they force any minimally rational miser to report previsions that obey the laws of probability. Second, they reveal the beliefs of expected utility maximizers because a miser who aims to maximize her *expected* payoff will invariably report a prevision for each proposition that coincides with her degree of belief for it. The fact that there exist scoring rules with these two properties is supposed to show that it is irrational to hold partial beliefs that violate the laws of probability.

Quadratic-loss functions ensure that rational previsions will be probabilities in virtue of

De Finetti's Lemma: In a prevision game scored by a quadratic-loss rule S , every prevision sequence p that violates the axioms of probability can be canonically associated with a sequence p^* that obeys the probability axioms and which *dominates* p in the sense that $S(p, \omega) > S(p^*, \omega)$ for all worlds ω .

In other words, for every sequence of previsions that violates the laws

actually exist, but they use them as a useful *idealization*. Insofar as a person is rational, it is claimed, she will pursue an abstract measure of overall satisfaction, *utility*, in the same way that a miser seeks wealth. The miser's craving for money is thus meant to mirror the universal desire for happiness.

of probability there is a sequence that obeys them whose penalty is strictly smaller *in every possible world*. No rational miser would ever choose to report previsions that are dominated in this way, since doing so would be tantamount to throwing away money.

I shall leave it to the reader to work out why the quadratic-loss rules penalize violations of Normalization and Non-negativity. For Additivity, imagine a person who reports previsions (0.6, 0.2) for (X, \sim X) when losses are given by the Brier score. This agent will incur a 10¢ penalty if X is true, and a 50¢ penalty if X is false. Figure 1 shows how she could have saved a sure penny by reporting the previsions (0.7, 0.3).




Figure 1. De Finetti's Lemma for $S((p, q), \omega) = 1/2[(\omega(X) - p)^2 + (\omega(\sim X) - q)^2]$. Previsions for (X, \sim X) appear as points in the (p, q)-plane. $V = \{(1,0), (0,1)\}$ is the set of all consistent truth-value assignments for X and \sim X. The line segment V^+ is V 's convex hull. It contains all (p, q) pairs with $p + q = 1$. Arc $C_1 = \{(p, q): S((p, q), 1) = 0.5\}$ is made up of points whose penalty is the same as that of (0.6, 0.2) when X is true. $C_0 = \{(p, q): S((p, q), 0) = 0.1\}$ contains all points whose penalty is the same as that of (0.6, 0.2) when X is false. The shaded region of dominance is the set of (p, q) pairs that have a smaller penalty than (0.6, 0.2) *whether X is true or false*. This region always intersects V^+ at (p^*, q^*) where $p^* = [p + (1 - q)]/2$ and $q^* = [p + (1 - q)]/2$. The Lemma says that one only has $(p, q) = (p^*, q^*)$ when $p + q = 1$.

This example mirrors the general case. If X is a finite sequence of propositions, then its consistent truth-value assignments form a family of binary sequences

$$V = \{(\omega(X_1), \omega(X_2), \dots, \omega(X_n)) : \omega \text{ a possible world}\}$$

within real n -dimensional space \mathfrak{R}^n . The *convex hull* of V is the subset V^+ of \mathfrak{R}^n whose points can be expressed as *weighted averages* of V 's elements. De Finetti showed that V^+ is the set of all prevision assignments for elements of X that obey the laws of probability. He then used the convexity of V^+ (the fact that it contains the line segment between any two of its points) to show that, for any quadratic-loss rule $S(\mathbf{p}, \omega) = \sum_i \lambda_i [\omega(X_i) - p_i]^2$ and any $\mathbf{p} \notin V^+$, there is a unique $\mathbf{p}^* \in V^+$ that minimizes $D(\mathbf{q}) = \sum_i \lambda_i [q_i - p_i]^2$ on V^+ and that this function has a lower S -score than \mathbf{p} does relative to *every* truth-value assignment in V .⁶

De Finetti's Lemma shows that a rational miser will always report previsions that obey the laws of probability when playing a prevision game scored by a quadratic-loss rule. But why think these previsions to have anything special to do with her degrees of belief? De Finetti often spoke as if there were no meaningful question to be asked here. A person's degrees of belief, he suggested, are *operationally defined* as *whatever previsions she would report in a game scored with a quadratic-loss rule*. This cannot be right. Aside from familiar difficulties with behaviorist interpretations of mental states, this view actually undermines itself. The problem is that it always makes sense to ask why a quadratic-loss function, rather than some other scoring rule, should be used to define degrees of belief. And, even if it is granted that a quadratic-loss rule should be used, one can still wonder whether all such rules will lead a rational miser to report the same previsions. After

6. Strictly speaking, this only establishes Additivity in the *finite* case. De Finetti did not go on to argue that the quadratic-loss rules enforce *countable* additivity because he felt a reasonable person should be able to assign the same, non-zero probability of winning to each ticket in a countably infinite lottery. As a number of authors have noted, however, de Finetti's argument for finite additivity extends easily to the infinite case. I have never seen a proof of this for the version of the Dutch book argument considered here. There are proofs for other versions (see Skyrms 1984, 21–23). Here is an (incomplete) sketch of how the proof would go: Normality and finite Additivity imply that any assignment \mathbf{p} of previsions to a countably infinite set of pairwise incompatible propositions $X = (X_1, X_2, X_3, \dots)$ is *square-convergent*, i.e. $\sum_i p_i^2$ is finite. V and V^+ are subsets of the space of square-convergent sequences. V^+ contains the countably additive prevision assignments for X . If we imagine previsions scored using a rule the quadratic $S(\mathbf{p}, \omega) = \sum_i \lambda_i (\omega(X_i) - p_i)^2$, then for any $\mathbf{p} \notin V^+$ and $\mathbf{q} \in V^+$ we can set $D(\mathbf{q}) = (\sum_i \lambda_i (q_i - p_i)^2)^{1/2}$ and minimize to find $\mathbf{p}^* \in V^+$. Calculation then shows that $S(\mathbf{p}^*, \omega) > S(\mathbf{p}, \omega)$ for all ω .

all, what prevents previsions from varying with changes in the weighting constants $\lambda_1, \dots, \lambda_n$? The point here is a general one. In the same way that it makes no sense to define “temperature” as “the quantity measured by thermometers” because it is impossible to know a priori either that such a quantity tracks any important physical property or that different thermometers will always assign similar values in similar circumstances, so too it makes no sense to define “degree of belief” as “the prevision reported in a quadratic-loss game” because it is impossible to know a priori either that previsions measure anything interesting or that different scoring rules elicit similar previsions in similar circumstances. It cannot be a definition which establishes that previsions reveal degrees of belief; it takes an argument.

As it turns out, de Finetti did not really need to rely on his operationism since he already had the required argument on hand (and indeed gave it). The reasoning turns on a substantive claim about the nature of practical rationality: viz., that a rational miser will always report previsions that maximize her *subjective expected utility*. She will, that is, always choose a prevision p_X for X that minimizes her *expected penalty* $Exp(p) = b(X)S(p, 1) + (1 - b(X))S(p, 0)$ where $b(X)$ is her degree of belief for X . It is not difficult to show that this function is uniquely minimized at $p_X = b(X)$ when S is any quadratic-loss function. This means that the previsions of *expected utility maximizers* do indeed reveal their degrees of belief. Since de Finetti’s Lemma shows that these previsions must obey the laws of probability, we are thus led to

The Dutch Book Theorem: If prudential rationality requires expected utility maximization, then any prudentially rational agent must have degrees of belief that conform to the laws of probability.

There are two main reasons why the Dutch book argument fails to convince people. First, there are some who reject the idea that prudential rationality requires expected utility maximization.⁷ I think these people are wrong, but will not argue the point here since for my purposes it is best to concede that the thesis is controversial so as to advertise the advantages of a defense of probabilism that does not presuppose it. A more significant problem has to do with the *pragmatic* character of the Dutch book argument. There is a distinction to be drawn between *prudential* reasons for believing, which have to do with the ways in which holding certain opinions can affect one’s happiness, and *epistemic* reasons for believing, which concern the *accuracy* of the opinions as representations of the world’s state. Since the Dutch book argument provides only a prudential rationale for conforming

7. The references here are too numerous to list. See Gardenfors and Shalin 1988.

one's partial beliefs to the laws of probability, it is an open question whether it holds any interest for epistemology. There are some who think it does not. Ralph Kennedy and Charles Chihara have written that:

The factors that are supposed to make it irrational to have a [probabilistically inconsistent] set of beliefs . . . are irrelevant, epistemologically, to the truth of the propositions in question. The fact (if it is a fact) that one will be bound to lose money unless one's degrees of belief [obey the laws of probability] just isn't epistemologically relevant to the *truth* of those beliefs. (1979, 30).

Roger Rosenkrantz has expressed similar sentiments, writing that the Dutch book theorem is a

roundabout way of exposing the irrationality of incoherent beliefs. What we need is an approach that . . . [shows] why incoherent beliefs are irrational from the perspective of the agent's purely *cognitive* goals. (1981, 214)

If this is right, then the pragmatic character of the Dutch book argument may well make it irrelevant to probabilism construed as a thesis in epistemology.

Proponents of the Dutch book argument might try to parry this objection by going pragmatist and denying that there is any sense in which the *epistemic* merits of a set of beliefs can outrun its prudential merits. Some old-line probabilists took this position, but it is unlikely to move anyone who feels the force of the Kennedy/Chihara/Rosenkrantz objection. There does seem to be a clear difference between appraising a system of beliefs in terms of the behavior it generates or in terms of its agreement with the facts. Unless the pragmatists can convincingly explain this intuition away it is hard to see how their view amounts to more than the bald assertion that there is no such subject as traditional epistemology. Probabilism is not worth that price.

More sophisticated probabilist responses acknowledge that partial beliefs can be criticized on nonpragmatic grounds, but they go on to suggest that imprudence, while not *constitutive* of epistemic failings, often *reliably indicates* them. People who choose means insufficient to their ends frequently do so because they weigh evidence incorrectly, draw hasty conclusions, engage in wishful thinking, or have beliefs that do not square with the facts. While this last flaw is no defect in rationality, it is reasonable to think that *systematic* deficiencies in practical reasoning *that do not depend on the truth or falsity of the reasoner's beliefs*, like the tendency of probabilistically inconsistent misers to throw away money, are symptoms of deeper flaws. If this is so, then

the Dutch book argument can be read as what Brian Skyrms (1984, 21–22) calls a “dramatic device” that provides a vivid pragmatic illustration of an essentially epistemic form of irrationality.

The kind of irrationality Skyrms has in mind is that of making *inconsistent value judgments*. As Ramsey first observed, an expected utility maximizer whose degrees of belief violate the axioms of probability cannot avoid assigning a utility to some prospect that is higher than the sum of the utilities she assigns to two others that together produce the same payoff as the first in *every* possible world. Her violations of the laws of probability thus leads her to commit both the prudential sin of squandering happiness and the epistemic sin of valuing prospects differently depending upon how they happen to be described. I want to agree that this is surely the right way to read the Dutch book argument: what the argument ultimately shows is that probabilistically inconsistent beliefs breed logically inconsistent preferences. The willingness to squander money is a side-effect of the more fundamental defect of having inconsistent desires. Still, even if we grant this point, it remains unclear why this should be counted an *epistemic* defect given that the inconsistency in question attaches to preferences or value judgments. It would be one thing if a Dutch book argument could show that the strengths of an agent’s *beliefs* vary with changes in the ways propositions happen to be expressed when she violates the laws of probability, but it cannot be made to show any such thing unless degrees of belief are assumed to obey the Additivity axiom from the start. The sort of inconsistency-in-valuing Skyrms decries is undeniably a serious shortcoming, but it remains unclear precisely what clearly irrational property of *beliefs* underlies it.⁸ In the end, the only way to answer the Chihara/Kennedy/Rosenkrantz objection is by presenting an argument that shows how having degrees of belief that violate the laws of probability engenders *epistemic* failings that go beyond their effects on an agent’s preferences.

3. The Concept of Gradational Accuracy. The main obstacle to such an argument is the lack of any compelling criterion of *epistemic success* for partial beliefs. Such a criterion has eluded probabilists because they have been slow to realize that full and partial beliefs “fit the facts” in different ways. The accuracies of full beliefs are evaluated on a *cate-*

8. One might be tempted here to say that it is the agent’s beliefs *about what is desirable* that are inconsistent. Aside from the fact that this would locate the epistemic flaw associated with my strongly believing both that it will be hot and that it will be cold tomorrow not in my beliefs about the weather but in my beliefs about the values of wagers, the underlying view that a desire can be understood as a kind of belief has serious difficulties. See Lewis 1988 and 1996 for relevant discussion.

gorical scale. The extent to which a full belief about X fits the facts is a matter of its “valence” (accept-X, reject-X, suspend belief), and X’s truth-value. Maximum (minimum) accuracy is attained when X is true (false) and accepted or when X is false (true) and rejected, and an intermediate value is obtained when belief is suspended. The “fit” between partial beliefs and the world is determined in a similar way except that, being attitudes that can come in a continuum of “valences,” their appropriate standard of accuracy must be a *gradational* one on which accuracy increases with the agent’s degrees of confidence in truths and decreases with her degrees of confidence in falsehoods.

To see what I have in mind, it is useful to consider Richard Jeffrey’s distinction between *guesses* and *estimates* of numerical quantities (Jeffrey 1986). When one tries to guess, say, the number of hits that a baseball player will get in his next ten at-bats, one aims to get the value exactly right. Guessing two hits when the batter gets three is just as wrong as guessing two hits when he gets ten. In guessing, closeness does not count. Not so for estimation. If the player gets five hits, it is better to have estimated that he would get three than to have estimated two or nine. Notice that, whereas it makes no sense to guess that a quantity will have a value that it cannot possibly have, it *can* make sense to estimate it to have such a value. One might, e.g., use a hitter’s batting average to estimate that he will get 3.27 hits in his next ten at-bats. Such an estimate can never be exactly right of course, but in estimation there is no special advantage to being *exactly* right; the goal is to get as *close* as possible to the value of the estimated quantity. In conditions of uncertainty it is often wise to “hedge one’s bets” by choosing a estimate that is sure to be off the mark by a little so as to avoid being off by a lot.

Following de Finetti, Jeffrey assumes that estimates must conform to the *laws of mathematical expectation*, and he identifies degrees of belief with *estimates of truth-values*. He is entirely right about the second point, but a bit too hasty with the first. When restricted to estimates of truth-values, the laws of mathematical expectation just are the laws of probability. Jeffrey takes this to provide a *justification* for requiring partial beliefs to satisfy the latter laws because he takes the former to be “as obvious as the laws of logic” (1986, 52). This, of course, is unlikely to convince anyone not already well disposed toward probabilism. The basic law of expectation is an additivity principle that requires a person’s expectation for a quantity to be the sum of her expectations of its summands, so that $Exp(F) = \sum_j Exp(F_j)$ when $F = \sum_j F_j$. No one who has qualms about additivity as it applies to degrees of belief is going to accept this stronger constraint without seeing a substantive argument.

The way to give a substantive argument, I believe, is to (a) grant Jeffrey's basic point that an agent's degree of belief for a proposition X is that number $b(X)$ that she is committed to using as her estimate of X 's truth-value when she recognizes that she will be evaluated for accuracy on a gradational standard appropriate for partial beliefs, and (b) argue that degrees of belief that obey the laws of probability are more accurate than those which do not when measured against this standard. What I have in mind here is a kind of "epistemic Dutch book argument" in which the relevant scoring rule assigns each credence function b and possible world ω a penalty $I(b, \omega)$ assessed in units of gradational inaccuracy. The rule I will gauge the extent to which the truth-value estimates sanctioned by b diverge from the truth-values that propositions would have were ω actual. My claim is going to be that, once we appreciate what I must look like, we will see that violations of the laws of probability always decrease the accuracy of partial beliefs.

Lest the reader think that I merely plan to restate the Dutch book argument and call it epistemology, let me highlight a crucial difference between my approach and that of de Finetti and Savage. Since a miser always aims to increase her fortune, de Finetti and Savage were at liberty to choose any scoring rule they wanted without having to worry about whether their subject would seek to minimize the penalties it assessed. This was advantageous for them because once they had discovered that the quadratic-loss rules rewarded the reporting of previsions that obey the laws of probability they could count on their subject to *want* to report such previsions. De Finetti and Savage did, of course, have to worry about whether their rules would induce a miser to report previsions that reveal her partial beliefs, which is why they needed to appeal to the principle of expected utility maximization. My problem is a mirror image of this. I cannot simply assume that my subjects will seek to minimize their penalties relative to any scoring rule I I might choose. The Norm of Gradational Accuracy portrays an epistemically rational agent as a kind of "accuracy miser." So, if a rule I does not measure gradational inaccuracy, then there is no good reason to think that such an agent will aim to minimize it. On the other hand, if I does measure gradational inaccuracy, then we can be sure that she will strive to have a system b of degrees of belief that minimizes $I(b, \omega_0)$ with respect to the actual world ω_0 . So, unless I can establish that my "scoring rule" really does measure inaccuracy in the epistemically relevant sense, I will have no grounds for concluding that we should care about its penalties. On the bright side, once I do find such a rule I can be sure that every epistemically rational agent will aim to have degrees of belief, not merely previsions, that minimize its values. This makes part

of my task easier than the one that faced de Finetti and Savage since I will not need to invoke any analogue of expected utility maximization.

To see why this is an advantage, consider a justification for probabilism offered by Roger Rosenkrantz (1981). While he does not invoke the distinction between categorical and gradational accuracy, it is not too much of a stretch to see Rosenkrantz asking the question that concerns us: assuming that the gradational *inaccuracy* of a system of degrees of belief can be measured by a function $I(\mathbf{b}, \omega)$, what properties must I have if it is going to be the sort of thing epistemically rational agents will seek to minimize. Rosenkrantz answers by introducing axioms that are meant to pick out the quadratic-loss rules as the only candidates for I . Among them we find:

Expected Accuracy Maximization: A rational agent should aim to hold a set of partial beliefs \mathbf{b} that minimizes her *expected inaccuracy*, i.e., for any partition X_1, X_2, \dots, X_n it must be true that $Exp(I(\mathbf{b}, \omega)) = \sum_i \mathbf{b}(X_i)I(\mathbf{b}, X_i) \geq Exp(I(\mathbf{b}^*, \omega)) = \sum_i \mathbf{b}(X_i)I(\mathbf{b}^*, X_i)$ for any alternative sets of degrees of belief \mathbf{b}^* .

Non-Distortion: The function $Exp(I(\mathbf{b}^*, \omega))$ attains a minimum at $\mathbf{b}(X_j) = \mathbf{b}^*(X_j)/\sum_i \mathbf{b}^*(X_i)$.

The quadratic-loss rules satisfy these conditions, and Rosenkrantz conjectures that they do so uniquely. While this may be so, the point is moot unless some non-circular rationale can be given for Expected Accuracy Maximization and Non-Distortion. Rosenkrantz does not offer any. Though I am happy to grant that both principles hold for partial beliefs that obey the axioms of probability, the problem is that they must also hold when the axioms are violated if they are to serve as premises in a justification for the fundamental dogma of probabilism. Here is a simple (but generalizable) example that shows why this cannot work. Let $\{X_1, X_2, X_3\}$ be a partition, and imagine someone with the probabilistically inconsistent beliefs $\mathbf{b}(X_1) = \mathbf{b}(X_2) = \mathbf{b}(X_3) = 1/3$ and $\mathbf{b}(X_2 \vee X_3) = 3/4$. If Rosenkrantz were right, this person would have to think that the most accurate degree of belief for X_1 is simultaneously $1/3 = \mathbf{b}(X_1)/[\mathbf{b}(X_1) + \mathbf{b}(X_2) + \mathbf{b}(X_3)]$ and $4/10 = \mathbf{b}(X_1)/[\mathbf{b}(X_1) + \mathbf{b}(X_2 \vee X_3)]$ because these are the answers that Non-distortion and Expected Accuracy Maximization sanction when applied to the partitions $\{X_1, X_2, X_3\}$ and $\{X_1, (X_2 \vee X_3)\}$ respectively. Perhaps Rosenkrantz would want to construe this inconsistency as an indication of irrationality, but unless he can offer us some independent rationale for his two principles we can just as well take the inconsistency to invalidate them as norms of epistemic rationality. The point here is basically the same as the one raised in connection with Jeffrey's

identification of estimates and expectations: we cannot hope to justify probabilism by assuming that rational agents should maximize the *expected* accuracy of their opinions because the concept of an expectation really only makes sense for agents whose partial beliefs already obey the laws of probability.

4. Measures of Gradational Accuracy. Despite this flaw in his argument, Rosenkrantz was right to think that a defense of the fundamental dogma should start from an analysis of inaccuracy measures, and that it should show that agents whose partial beliefs violate the axioms of probability are always less accurate than they need to be. I will provide a defense along these lines by formulating and justifying a set of constraints on measures of gradational *inaccuracy*, and then showing that any function that meets these constraints will encourage conformity to the laws of probability in the strongest possible manner. It will turn out that, relative to any such measure, a system of partial beliefs that violates the axioms of probability can always be replaced by a system that both obeys the axioms and better fits the facts no matter what the facts turn out to be.

In developing these ideas, I will speak as if gradational accuracy can be precisely quantified. This may be unrealistic since the concept of accuracy for partial beliefs may simply be too *vague* to admit of sharp numerical quantification. Even if this is so, however, it is still useful to pretend that it can be so characterized since this lets us take a “supervaluationist” approach to its vagueness. The supervaluationist idea is that one can understand a vague concept by looking at all the ways in which it can be made precise, and treating facts about the properties that all its “precisifications” share as facts about the concept itself. In this context a “precisification” is a real function that assigns a definite inaccuracy score $I(\mathbf{b}, \omega)$ to each set of degrees of belief \mathbf{b} and world ω . In what follows, I am going to be interested not so much in what the function I is, but in the properties that all reasonable “precisified” measures of gradational inaccuracy must share.

Let me begin by codifying the notation. The measure I is defined over pairs in $\mathbf{B} \times \mathcal{V}$, where \mathbf{B} is the family of all credence functions defined on a countable⁹ Boolean algebra of propositions Ω and \mathcal{V} is the subset of \mathbf{B} containing all consistent truth-value assignments to members of Ω . We will continue referring to these truth-value assign-

9. It does no harm to assume that Ω is *countable* since violations of the laws of probability always occur in countable sets. On an uncountable algebra of propositions the probabilist requirement is that degrees of belief should obey the probability axioms on every *countable subalgebra*.

ments as “possible worlds” and using “ ω ” as a generic symbol for them. The collection of all probability functions in \mathbf{B} is V ’s convex hull V^+ . $\mathbf{B} \sim V^+$ is thus the set of all assignments of degrees of belief to the propositions in Ω that violate the laws of probability. The set \mathbf{B} is endowed with a great deal of geometrical structure. It always contains a unique “line” $\mathbf{L} = \{\lambda\mathbf{b} + (1-\lambda)\mathbf{b}^*: \lambda \in \mathfrak{R}\}$ that passes through any two of its “points” \mathbf{b} and \mathbf{b}^* . The *line segment* from \mathbf{b} to \mathbf{b}^* , hereafter \mathbf{bb}^* , is the subset of \mathbf{L} for which λ falls between zero and one. A function $[\lambda\mathbf{b} + (1-\lambda)\mathbf{b}^*]$ that falls on this segment is called a *mixture* of \mathbf{b} and \mathbf{b}^* since it assigns each $X \in \Omega$ a “mixed” value of $\lambda\mathbf{b}(X) + (1-\lambda)\mathbf{b}^*(X)$. This *mixture* effects a kind of compromise between \mathbf{b} and \mathbf{b}^* when the two differ. If $\lambda > 1/2$ the compromise favors the \mathbf{b} beliefs since $\lambda\mathbf{b}(X) + (1-\lambda)\mathbf{b}^*(X)$ is always closer to $\mathbf{b}(X)$ than to $\mathbf{b}^*(X)$. The reverse occurs when $\lambda < 1/2$. The *even mixture* ($\lambda = 1/2$) is a “fair” compromise that sets X ’s degree of belief exactly halfway between $\mathbf{b}(X)$ and $\mathbf{b}^*(X)$. A number of the constraints to be imposed below will exploit this geometry of lines and segments.

Our first axiom says that inaccuracy should be non-negative, that small changes in degrees of belief should not engender large changes in accuracy, and that inaccuracy should increase without limit as degrees of belief move further and further from the truth-values of the propositions believed.

Structure: For each $\omega \in V$, $I(\mathbf{b}, \omega)$ is a non-negative, *continuous* function of \mathbf{b} that goes to infinity in the limit as $\mathbf{b}(X)$ goes to infinity for any $X \in \Omega$.

This weak requirement should be uncontroversial given that gradational accuracy is supposed to be a matter of “closeness to the truth.”

Our next constraint stipulates that the “facts” which a person’s partial beliefs must “fit” are exhausted by the truth-values of the propositions believed, and that the only aspect of her opinions that matter is their strengths.

Extensionality: At each possible world ω , $I(\mathbf{b}, \omega)$ is a function of nothing other than the truth-values that ω assigns to propositions in Ω and the degrees of confidence that \mathbf{b} assigns these propositions.

Most objections to Extensionality conflate the task of finding a measure of accuracy for partial beliefs with the more ambitious project of defining an *epistemic utility* function that gauges the overall goodness of a system of partial beliefs in all epistemologically relevant respects.¹⁰ Accuracy is only one virtue among many that we want our opinions to

10. For an excellent recent discussion of epistemic utility, see Maher 1993.

possess. Ideally, a person will hold beliefs that are informative, simple, internally coherent, well-justified, and connected by secure causal links to the world. A notion of epistemic utility will balance off all these competing desiderata to provide an “all-in” measure of doxastic quality. While accuracy will be a strongly-weighted factor in any such measure, it will not be the only factor. Since properties like the informativeness of a belief or its degree of justification are not extensional, epistemic utility cannot be either. Extensionality does make sense for gradational accuracy, however, since gradational accuracy is supposed to be the analogue of *truth* for partial beliefs. Just as the accuracy of a full belief is a function of its attitudinal “valence” (accept/reject/suspend judgment) and its truth-value, so too the accuracy of a partial belief should be a function of its “valence” (degree) and truth-value.

A second objection to Extensionality is that it does not take *verisimilitude* into account.¹¹ Here is how the complaint might go:

Copernicus (let us suppose) was exactly as confident that the earth’s orbit is circular as Kepler was that it is elliptical. However, both were wrong since the gravitational attraction of the moon and the other planets causes the earth to deviate slightly from its largely elliptical path. Extensionality rates the two thinkers as *equally* inaccurate since both believed a falsehood to the same high degree. Still Kepler was obviously nearer the mark, which suggests that evaluations of accuracy must be sensitive not only to the truth-values of the propositions involved, but also to how *close* false propositions come to being true.

I am happy to admit that Kepler held more accurate beliefs than Copernicus did, but I think the sense in which they were more accurate is best captured by an extensional notion. While Extensionality rates Kepler and Copernicus as equally inaccurate when their false beliefs about the earth’s orbit are considered *apart from their effects on other beliefs*, the advantage of Kepler’s belief has to do with the other opinions it supports. An agent who strongly believes that the earth’s orbit is elliptical will also strongly believe many more *truths* than a person who believes that it is circular (e.g., that the average distance from the earth to the sun is different in different seasons). This means that the overall effect of Kepler’s inaccurate belief was to improve the *extensional* accuracy of his system of beliefs as a whole. Indeed, this is why his theory won the day. I suspect that most intuitions about falsehoods being “close to the truth” can be explained in this way, and that they therefore pose no real threat to Extensionality.

11. Thanks to Bob Batterman for helping me think this issue through.

Our third axiom requires the accuracy of a system of degrees of belief to be an *increasing* function of the believer's degree of confidence in any truth and a decreasing function of her degree of confidence in any falsehood.

Dominance: If $b(Y) = b^*(Y)$ for every $Y \in \Omega$ other than X , then $I(b, \omega) > I(b^*, \omega)$ if and only if $|\omega(X) - b(X)| > |\omega(X) - b^*(X)|$.

This principle really says two things. First, it lets us speak of the accuracy of each *individual* degree of belief taken in isolation from the belief system as a whole. Second, it says that the accuracy of $b(X)$ always increases as it approaches $\omega(X)$. Thus, moving one's degree of belief for X closer to X 's truth-value improves accuracy *no matter what one's other degrees of belief might be*. Were this not the case one could have a perverse incentive to *lower* one's degree of belief in a proposition for whose truth one has strong evidence because doing so would increase overall accuracy.

To see how bizarre these incentives can be, consider the *calibration index*, a measure of accuracy for degrees of belief that Bas van Fraassen and Abner Shimony have each tried to use in a vindication of probabilism similar to the one sought here. As Wesley Salmon (1988) noted, many probabilists are attracted to *frequency driven* accounts of subjective probability. The *truth-frequency* of a family of propositions $X = \{X_1, X_2, \dots, X_n\}$ at a world ω is the proportion of the X_i that hold in ω , so that $\text{Freq}(X, \omega) = [\omega(X_1) + \omega(X_2) + \dots + \omega(X_n)]/n$. It is easy to show that an agent who has well-defined degrees of belief for all X 's elements can only satisfy the axioms of probability if her *expected* frequency of truths in X is equal to her *average* degree of belief for the various X_i , so that $\text{Exp}(\text{Freq}(X)) = [b(X_1) + \dots + b(X_n)]/n$. A special case of this is

The Calibration Theorem: If an agent assigns the same degree of belief x to every proposition in X , then a necessary condition for her degrees of belief to satisfy the axioms of probability is that her expectation for the frequency of truths in X must be x .

This seems to get at something deep about partial beliefs. What can it mean, after all, to assign degree of belief x to X if not to think something like, "Propositions like X are true about x proportion of the time"? Moreover, unlike the principle of mathematical expectation from which it follows, the Calibration Theorem does not presuppose probabilism in any obvious way. Perhaps the thing to do is to replace "satisfy the axioms of probability" by "be rational" and "expectation" by "estimate," and to treat the Calibration Theorem as a conceptual truth about degrees of belief. And, if one does so, the accuracy of a set

of degrees of belief can be analyzed as a function of the discrepancy between the relative frequency estimates it sanctions and the actual relative frequencies.

The meteorologist A. Murphy found a way to measure this discrepancy (Murphy 1973). For any credence function \mathbf{b} defined over a *finite* family of propositions X , one can always subdivide X into disjoint *reference classes* $X_j = \{X \in X: \mathbf{b}(X) = b_j\}$, where $\{b_1, \dots, b_n\}$ lists all the values that \mathbf{b} assumes on X . The Calibration Theorem tells us that b_j is the only estimate for $\mathbf{Freq}(X_j)$ that \mathbf{b} can sanction. Murphy characterized the divergence of these estimates from the actual frequencies at world ω using a quantity called the *calibration index* $\mathcal{Cal}(\mathbf{b}, X, \omega) = \sum_j (n_j/n)[\mathbf{Freq}(\omega(X_j)) - b_j]^2$ where n is the number of propositions in X and n_j is the number of propositions in X_j . The function \mathbf{b} is *perfectly calibrated* when $\mathcal{Cal}(\mathbf{b}, X, \omega) = 0$. In this case, half the elements of X assigned value 1/2 are true, two-fifths of those assigned value 2/5 are true, three-fourths of those assigned value 3/4 are true, and so on.

Some have championed calibration as the best measure of “fit” between partial beliefs and the world. Van Fraassen, for example, has written that calibration “plays the conceptual role that truth . . . has in other contexts” (1983, 301), and has suggested that the appropriate analogue of consistency for degrees of belief is *calibrability*, the ability to be embedded within ever richer systems of beliefs whose calibration scores can be made arbitrarily small. He and Abner Shimony (1988) have even sought to vindicate probabilism by arguing, in different ways, that the only way to achieve calibrability with respect to finite sets of propositions is by having degrees of belief that conform to the laws of probability. If either of these arguments had succeeded we would have had our nonpragmatic vindication of probabilism.

They fail for two reasons. First, van Fraassen and Shimony need to employ very strong structural assumptions that are not well motivated as requirements of rationality. While the two assumptions are similar, van Fraassen’s is easier to state because he deals only with propositions of the monadic form “ x is A .” He requires that for any assignment \mathbf{b} of degrees of belief to the elements of a set X of such propositions it should be possible to extend \mathbf{b} to a function \mathbf{b}^* defined on a superset X^* of X in such a way that each proposition “ x is A ” in X can be associated with a subset in X^* of the form

$$X(x, A) = \{x \text{ is } A, x_1 \text{ is } A, x_2 \text{ is } A, \dots, x_k \text{ is } A\}$$

where (a) k may be any positive integer, (b) $\mathbf{b}^*(x_j \text{ is } A) = \mathbf{b}(x \text{ is } A)$ for every j , and (c) the propositions in $X(x, A)$ are *logically independent* of one another. In effect, van Fraassen is introducing dummy propo-

sitions to ensure that each element of X can be embedded in a probabilistically homogenous reference class of any chosen truth-frequency. Shimony uses a somewhat more general condition, his E_1 (1988, 156–157), to achieve substantially the same end. These are extremely strong, and rather ad hoc, assumptions, and it is not at all surprising that grand conclusions can be deduced from them. What remains unclear, however, is why rational degrees of belief should be required to satisfy any such conditions.

But, even supposing that it is possible to show that they should, a more substantive problem with the van Fraassen/Shimony approach is that calibration is simply not a reasonable measure of accuracy for partial beliefs.¹² Consider the following table, which gives four sets of degrees of belief for propositions in $X = \{X_1, X_2, X_3, X_4\}$ and their calibration scores at a world ω in which X_1 and X_2 are true and X_3 and X_4 are false:

	b_1	b_2	b_3	$X_j(\omega)$
X_1	1/2	1	9/10	1
X_2	1/2	1	9/10	1
X_3	1/2	1/10	1/2	0
X_4	1/2	0	1/2	0
<i>Cal</i>	0	1/400	13/100	0

Figure 2. Calibration Scores.

Notice that b_1 is *better* calibrated than b_2 even though *all* of b_2 's values are *closer* to the actual truth-values than those of b_1 . This happens because each individual degree of belief can affect the overall calibration of its credence function not only by being closer to the truth-value of the proposition believed, but by manipulating the family of subsets relative to which calibration is calculated. To see why this is a problem imagine that an agent with degrees of belief b_3 who has strong evidence for X_1 and X_2 , somehow learns that exactly two of the X_j hold, without being told which ones. What should he do with this information? One might think that a rational believer would lower his estimates for X_3 and X_4 to nearly zero and keep his estimates for X_1 and X_2 close to one. If we equate accuracy with good calibration, however, this is wrong! The best way for our agent to improve his calibration score (indeed to ensure that it will be zero) is to keep his estimates for X_3

12. My discussion here is indebted to Seidenfeld 1985.

and X_4 fixed, ignore all his evidence, and lower his estimates for X_1 and X_2 to $1/2$. The Dominance requirement rules out this sort of absurdity.

Our fourth axiom says that differences among possible worlds that are not reflected in differences among truth-values of proposition that the agent believes should have no effect on the way in which accuracy is measured.

Normality: If $|\omega(X) - \mathbf{b}(X)| = |\omega^*(X) - \mathbf{b}^*(X)|$ for all $X \in \Omega$, then $I(\mathbf{b}, \omega) = I(\mathbf{b}^*, \omega^*)$.

In the presence of the other conditions, this merely says that the standard of gradational accuracy must not vary with changes in the world's state that do not effect the truth-values of believed propositions. Were this not so there would be no uniform notion of "what it takes" for a system of partial beliefs to fit the facts.

Our final two constraints concern mixtures of credence functions.

Weak Convexity: Let $\mathbf{m} = (1/2\mathbf{b} + 1/2\mathbf{b}^*)$ be the midpoint of the line segment between \mathbf{b} and \mathbf{b}^* . If $I(\mathbf{b}, \omega) = I(\mathbf{b}^*, \omega)$, then it will always be the case that $I(\mathbf{b}, \omega) \geq I(\mathbf{m}, \omega)$ with identity only if $\mathbf{b} = \mathbf{b}^*$.

Symmetry: If $I(\mathbf{b}, \omega) = I(\mathbf{b}^*, \omega)$, then for any $\lambda \in [0,1]$ one has $I(\lambda\mathbf{b} + (1-\lambda)\mathbf{b}^*, \omega) = I((1-\lambda)\mathbf{b} + \lambda\mathbf{b}^*, \omega)$.

To see why Weak Convexity is a reasonable constraint on gradational inaccuracy notice that in moving from \mathbf{b} to \mathbf{m} an agent would alter each of degree of belief $\mathbf{b}(X)$ by adding an increment of $\mathbf{k}(X) = 1/2[\mathbf{b}^*(X) - \mathbf{b}(X)]$. She would add the same increment of $\mathbf{k}(X)$ to each $\mathbf{m}(X)$ in moving from \mathbf{m} to \mathbf{b}^* . To put it in geometrical terms, the "vector" \mathbf{k} that she must add to \mathbf{b} to get \mathbf{m} is the same as the vector she must add to \mathbf{m} to get \mathbf{b}^* . Furthermore, since $\mathbf{b}^* = \mathbf{b} + 2\mathbf{k}$ the change in belief involved in going from \mathbf{b} to \mathbf{b}^* has the *same direction* but a *doubly greater magnitude* than change involved in going from \mathbf{b} to \mathbf{m} . This means that the former change is more *extreme* than the latter in the sense that, for *every* proposition X , both changes alter the agent's degree of belief for X in the same direction, either by moving it closer to one or closer to zero, but the \mathbf{b} to \mathbf{b}^* change will always move $\mathbf{b}(X)$ twice as far as the \mathbf{b} to \mathbf{m} change moves it. Weak Convexity is motivated by the intuition that extremism in the pursuit of accuracy is no virtue. It says that if a certain change in a person's degrees of belief does not improve accuracy then a more radical change in the same direction and of the same magnitude should not improve accuracy either. Indeed, this is just what the principle says. If it did not hold, one could have absurdities like this: "I raised my confidence levels in X and Y and my beliefs became less accurate overall, so I raised my

confidence levels in X and Y again, by exactly the same amounts, and the initial accuracy was restored.”

To understand the rationale for Symmetry observe first that, when \mathbf{b} and \mathbf{b}^* are equally accurate at ω , Weak Convexity entails that there will always be a unique point on the interior of the line segment between them that minimizes inaccuracy over the segment, i.e., there will be a $\mathbf{c} = \mu\mathbf{b} + (1-\mu)\mathbf{b}^*$ with $0 < \mu < 1$ such that $I(\lambda\mathbf{b} + (1-\lambda)\mathbf{b}^*, \omega) \geq I(\mathbf{c}, \omega)$ for all λ with $0 \leq \lambda \leq 1$.¹³ If \mathbf{c} were not the midpoint of $\mathbf{b}\mathbf{b}^*$, then it would have to be closer to \mathbf{b} or to \mathbf{b}^* . Given the initial symmetry of the situation this would amount to an unmotivated bias in favor of one set of beliefs or the other. If $\mathbf{c} = 1/4\mathbf{b} + 3/4\mathbf{b}^*$, for example, then \mathbf{c} would lie between \mathbf{b}^* and the midpoint of $\mathbf{b}\mathbf{b}^*$. This would mean that a person who held the \mathbf{b} beliefs would need to alter her opinions more radically than a person who held the \mathbf{b}^* beliefs in order to attain the maximum accuracy along $\mathbf{b}\mathbf{b}^*$. The reverse would be true if $\mathbf{c} = 3/4\mathbf{b} + 1/4\mathbf{b}^*$. Symmetry rules this sort of thing out. It says that when \mathbf{b} and \mathbf{b}^* are equally accurate there can be no grounds, based on considerations of accuracy alone, for preferring a “compromise” that favors \mathbf{b} to a symmetrical compromise that favors \mathbf{b}^* . It does this by requiring that the change in belief that moves an agent a proportion λ along the line segment from \mathbf{b} toward \mathbf{b}^* has the same overall effect on her accuracy as a “mirror image” change that moves her the same proportion λ along the line segment from \mathbf{b}^* toward \mathbf{b} .

Structure, Extensionality, Normality, Dominance, Weak Convexity, and Symmetry are the only constraints on measures of gradational accuracy we need to vindicate the fundamental dogma of probabilism. Those who find these conditions compelling, and who agree with my analysis of partial beliefs as estimates of truth-value, are thereby committed to thinking that epistemically rational degrees of belief must obey the laws of probability. Those who deny this will either need to explain where my conditions go wrong or will have to dispute my analysis of partial beliefs. For the reasons presented, I do not believe either line of attack will succeed.

5. Vindicating the “Fundamental Dogma”. In this section we will see how any system of degrees of belief that violates the axioms of probability can be replaced by a system that both obeys these axioms and is more accurate relative any assignment of truth-values to the propositions believed. The aim is to prove the

Main Theorem: If gradational inaccuracy is measured by a func-

13. The proof of this fact is essentially the same as the proof of Lemma-1, below.

tion I that satisfies Structure, Extensionality, Normality, Dominance, Weak Convexity, and Symmetry, then for each $c \in B \sim V^+$ there is a $c^* \in V^+$ such that $I(c, \omega) > I(c^*, \omega)$ for every $\omega \in V$.

Begin the proof by defining a map $D(b, c) = I(\omega + (b - c), \omega)$ where $\omega + (b - c)$ is defined by $(\omega + b - c)(X) = \omega(X) + b(X) - c(X)$. (I have chosen the symbol “ D ” here to suggest the notion of a distance function.)

The following facts are simple consequences of the conditions we have imposed on I : (Proofs are left to interested readers, but the axioms needed for each case are given.)

- I. $D(\bullet, c)$ is continuous for each $c \in B$. [Structure]
- II. D 's value does not depend on the choice of $\omega \in V$. [Structure]
- III. $D(b, c)$ goes to infinity as $b(X)$ goes to infinity for any $X \in \Omega$. [Structure]
- IV. $D(b, c) \geq D(b^*, c^*)$ if $|b(X) - c(X)| \geq |b^*(X) - c^*(X)|$ holds for all $X \in \Omega$, and the former inequality is strict if the latter is strict for some X . [Dominance]
- V. If c^* lies on the line segment bc and if $c^* \neq b$, then $D(b, c) > D(c^*, c)$. [via IV]
- VI. $D(b, c) = D(b^*, c)$ if and only if $D(\bullet, c)$ has a unique minimum along the line segment bb^* at its midpoint $1/2b + 1/2b^*$. [Symmetry, Weak Convexity]

We will use these facts to prove a series of lemmas that establish the Main Theorem.

Let c be any fixed element of $B \sim V^+$. Our first lemma shows how to select c^* , the point in V^+ that is “closer to the truth” than c is no matter what the truth turns out to be.

LEMMA-1: There is a point $c^* \in V^+$ such that the function $D(\bullet, c)$ attains its *unique* minimum on V^+ at c^* .

PROOF: A classic result from point-set topology says that a continuous, real-valued function defined on a closed, bounded region always attains a minimum on that region. Since V^+ is closed and bounded it follows from (I) that there is a point $c^* \in V^+$ with $D(c^*, c) \leq D(b, c)$ for all $b \in V^+$. To see why this minimum is unique, suppose it is attained by another $b^* \in V^+$. Since $D(b^*, c) = D(c^*, c)$, fact (VI) entails that $D(\bullet, c)$ assumes a unique minimum on the line segment c^*b^* at its midpoint $1/2c^* + 1/2b^*$. Since V^+ is convex it will contain this midpoint, which contradicts the hypothesis that c^* minimizes $D(\bullet, c)$ on V^+ . Q.E.D.

Given Lemma-1, we can prove Main Theorem by showing that $I(c, \omega)$

$> I(c^*, \omega)$ for all $\omega \in V$. Start by selecting an arbitrary ω . We may assume that c^* and ω are distinct, and thus that $D(c^*, c) < D(\omega, c)$, since the desired inequality follows trivially from (IV) if they are identical. Let $L = \{\lambda c^* + (1-\lambda)\omega : \lambda \in \mathfrak{R}\}$ be the line in B that contains c^* and ω , and let $R = \{\lambda c^* + (1-\lambda)\omega : \lambda \geq 1\}$ be the ray of L that begins at c^* but does not contain ω .

LEMMA-2: There is a point m on R such that (a) m uniquely minimizes $D(\bullet, c)$ on R , (b) c^* is an element of the segment of L that runs between m and ω , and (c) $I(m, \omega) \geq I(c^*, \omega)$.

PROOF: Fact (III) entails that $D(\bullet, c)$ goes to infinity on R as λ does. Given that $D(c^*, c) < D(\omega, c)$ it follows from (I), and the Intermediate Value Theorem, that there is a point k on R such that $D(k, c) = D(\omega, c)$. Let $m = 1/2k + 1/2\omega$ be the midpoint of the line segment $k\omega$. By (VI), m is the unique minimum of $D(\bullet, c)$ on this segment. m cannot lie strictly between c^* and ω on L because it would then be contained in V^+ , which would entail that c^* does not minimize $D(\bullet, c)$ on V^+ . Thus, c^* must be on segment $m\omega$, and (V) entails that $I(m, \omega) \geq I(c^*, \omega)$, with the equality strict if $c^* \neq m$. Q.E.D.

Given these two Lemmas, the Main Theorem follows if it can be shown that $I(c, \omega) > I(m, \omega)$. This is one of those cases where a picture is worth a thousand words.

LEMMA-3: $I(c, \omega) > I(m, \omega)$.

PROOF: By the construction of Lemma-2 we know that $D(k, c) = D(\omega, c)$. Since c minimizes $D(\bullet, c)$ on the line segment from k to $2c - k$, (VI) entails that $D(k, c) = D(2c - k, c)$. Together these identities yield

$$(A) \quad D(\omega, c) = D(2c - k, c).$$

Given (A), fact (VI) entails that $D(\bullet, c)$ attains a unique minimum on line segment between ω and $2c - k$ at $[1/2(\omega - k) + c]$. It follows that




Figure 3. The Key Lemma in the Proof of the Main Theorem: $d_1 > d_2$.

$$(B) \quad D(\omega, c) > D(1/2(\omega - k) + c, c).$$

Since D is a symmetric function of its two arguments this means that

$$(C) \quad D(c, \omega) > D(c, 1/2(\omega - k) + c).$$

We can now use the definition of D to obtain

$$\begin{aligned} I(c, \omega) &= D(c, \omega) > D(c, 1/2(\omega - k) + c) \\ &= I(\omega + (c - [1/2(\omega - k) + c]), \omega) \\ &= I(1/2\omega + 1/2k, \omega) \\ &= I(m, \omega). \end{aligned}$$

So, we have shown that $I(c, \omega) > I(m, \omega)$. Q.E.D.

Since we already know from Lemma-2 that $I(m, \omega) \geq I(c^*, \omega)$, we obtain the inequality $I(c, \omega) > I(c^*, \omega)$ from Lemma-3. This completes the proof of the Main Theorem. It is thus established that degrees of belief that violate the laws of probability are invariably less accurate than they could be. Given that an epistemically rational agent will always strive to hold partial beliefs that are as accurate as possible, this vindicates the fundamental dogma of probabilism.

6. Some Loose Ends. The foregoing results suggests two further lines of investigation. First, it would be useful to know what functions obey the constraints imposed on I . Second, to apply the Main Theorem in realistic cases we need to understand how it applies to partial beliefs that do not admit of measurement in precise numerical degrees.

I cannot now specify the class of functions that satisfy my axioms, but I do know it is not empty. The quadratic-loss rules are among its elements, as is any map $I(\mathbf{b}, \omega) = F(\sum_{X \in \Omega} \lambda_X [\omega(X) - \mathbf{b}(X)]^2)$ where F is a continuous, strictly increasing real function. The proofs of these claims are, however, beyond the scope of this paper. I am not certain whether there are other functions that meet the requirements,¹⁴ but I suspect there are.

Turning to the second issue, the Main Theorem tells us that partial beliefs whose strengths can be measured in precise numerical degrees must conform to the laws of probability, but its import is less clear for partial beliefs specified in more realistic ways. Most probabilists recognize that opinions are often too vague to be pinned down in numerical terms, and it has therefore become standard to represent a person's partial beliefs not by some single credence function but by the class of all credence functions consistent with her opinions. One then thinks of a doxastic state not as a single element of \mathbf{B} but as one of its subsets \mathbf{B}^* .

14. One large class of functions that do not satisfy them (because they violate Symmetry) are the ℓ_p -norms: $I(\mathbf{b}, \omega) = (\sum_{X \in \Omega} \lambda_X [\omega(X_i) - \mathbf{b}(X_i)]^p)^{1/p}$, for $p \geq 1$ other than $p = 2$.

The most minimal probabilistic consistency requirement for partial beliefs that are modeled in this way is that there should be *at least one* probability among the elements of B^* . In other words, an epistemically rational agent's partial beliefs should always be *extendible* to some system of degrees of belief that satisfy the axioms of probability. The Main Theorem provides a compelling rationale for this requirement because if B^* contained no probabilities then *every* way of making the agent's opinions precise would result in a system of degrees of belief that are less accurate than they could otherwise be. It would then be determinately the case that the agent's partial beliefs are not as accurate as they could be because every precisification of them would yield a credence function that is less accurate than it could be.

One of the best things about looking at matters in this way is that it helps to make sense of some old results pertaining to the probabilistic representation of ordinal confidence rankings. In a seminal paper, Kraft, et al. (1959) presented a set of necessary and sufficient conditions for a *comparative probability ranking* to be represented by a probability. We may think of such a ranking as a pair of relations ($.>., .\geq.$) defined on Ω , where $X .>. Y$ and $X .\geq. Y$ mean, respectively, that the agent is more confident in X than in Y , or as confident in X as in Y . The conditions Kraft et al. laid down can be expressed in a variety of ways, but the most tractable formulation is due to Dana Scott (1964). Say that two *ordered sequences* of (not necessarily distinct) propositions (X_1, X_2, \dots, X_n) and (Y_1, Y_2, \dots, Y_m) drawn from Ω are *isovalent* (my term) when the number of truths that appear in the first is necessarily identical to the number that appear in the second, so that $\omega(X_1) + \omega(X_2) + \dots + \omega(X_n) = \omega(Y_1) + \omega(Y_2) + \dots + \omega(Y_m)$ holds at *every* world ω . The important thing about isovalence is that a probability function β will always be additive over isovalent sequences, so that $\sum_i \beta(X_i) = \sum_i \beta(Y_i)$ when (X_1, X_2, \dots, X_n) and (Y_1, Y_2, \dots, Y_m) are isovalent. Scott introduced the following constraint on confidence rankings to ensure that all their representations would have this generalized additive property:

Scott's Axiom: If (X_1, X_2, \dots, X_n) and (Y_1, Y_2, \dots, Y_m) are isovalent, it should never be the case that $X_i .\geq. Y_i$ for every $i = 1, 2, \dots, n$ where $X_j .>. Y_j$ for some j .

He then proved that, for finite Ω , Scott's Axiom (plus a nontriviality requirement) is necessary and sufficient for the existence of a probability representation for ($.>., .\geq.$).

Commentators have not known what to make of Scott's condition. Scott himself worried about its "non-Boolean" nature. Terrence Fine points out, quite rightly, that it makes essential reference to *sums* of

propositions which generally will not be propositions themselves. A reasonable theory of comparative probability, he writes, should be, “concerned only with [propositions]. Why should we be concerned about objects that have no reasonable interpretation in terms of random phenomena?” (1973, 24) Peter Forrest, commenting on a condition of his own that is equivalent to Scott’s Axiom, writes:

My results are largely negative, I motivate the search for a certain kind of representation and I provide a condition which, given various intuitive rationality constraints, is necessary, sufficient and non-redundant. Unfortunately, this condition is not itself an intuitive rationality constraint. That is why my results are negative. Their chief purpose is to throw out a challenge. Is it possible to provide an intuitive rationality constraint that implies [Scott’s Axiom]? (1989, 280)

Fortunately, we already have one! Scott’s Axiom is just the requirement one would impose if one wanted partial beliefs to be gradationally accurate. If (X_1, X_2, \dots, X_n) and (Y_1, Y_2, \dots, Y_m) are isovalent, then every logically consistent set of truth-value assignments ω will be found somewhere in the bounded, closed, convex set

$$U = \{b \in \mathbf{B}: b(X_1) + \dots + b(X_n) = b(Y_1) + \dots + b(Y_m), \text{ for } 0 \leq b(X_i), b(Y_j) \leq 1\}.$$

If $X_i \succ Y_i$ for all i with $X_j \succeq Y_j$ for some j , then any credence function c that represents these beliefs will satisfy $[c(X_1) + \dots + c(X_n)] > [c(Y_1) + \dots + c(Y_m)]$, which means that c will lie *outside* U . By recapitulating our argument for the Main Theorem we can find a point $c^* \in U$ such that $I(c, \omega) > I(c^*, \omega)$ for *every* world ω . Thus, once we start thinking in terms of gradational accuracy, Scott’s Axiom can be interpreted as a constraint that prevents people from having partial beliefs that are less accurate than they need to be. This, as we have seen, is something to be avoided on pain of *epistemic* irrationality.

REFERENCES

- Anscombe, G. E. M. (1957), *Intention*. Oxford: Basil Blackwell.
 Brier, G. (1950) “Verification of Forecasts Expressed in Terms of Probability”, *Monthly Weather Review* 78: 1–3.
 Chisholm, R. (1977), *Theory of Knowledge*, 2nd ed. New York: Prentice Hall.
 de Finetti, B. (1974), *Theory of Probability*, vol. 1. New York: John Wiley and Sons.
 Fine, T. (1973), *Theories of Probability*. New York: Academic Press.
 Foley, R. (1987), *The Theory of Epistemic Rationality*. Cambridge, MA: Harvard University Press.
 Forrest, P. (1989), “The Problem of Representing Incompletely Ordered Doxastic Systems”, *Synthese* 79: 18–33.

- Gärdenfors, P. and P. Shalin (1988), *Decision, Probability, and Utility*. Cambridge: Cambridge University Press.
- Jeffrey, R. (1986), "Probabilism and Induction", *Topoi* 5: 51–58.
- . (1992), "Probability and the Art of Judgment", in *Probability and the Art of Judgment*. Cambridge: Cambridge University Press, 44–76.
- Kennedy, R. and C. Chihara (1979) "The Dutch Book Argument: its Logical Flaws, its Subjective Sources", *Philosophical Studies* 36: 19–33.
- Kraft, C., J. Pratt, and A. Seidenberg (1959), "Intuitive Probability on Finite Sets", *Annals of Mathematical Statistics* 30: 408–419.
- Lewis, David (1988), "Desire as Belief", *Mind* 97: 323–332.
- . (1996), "Desire as Belief II", *Mind* 105: 303–313.
- Maher, P. (1993), *Betting on Theories*. Cambridge: Cambridge University Press.
- Murphy, A. (1973), "A New Vector Partition of the Probability Score", *Journal of Applied Meteorology* 12: 595–600.
- Ramsey, F. P. (1931), "Truth and Probability", in R. B. Braithwaite (ed.), *The Foundations of Mathematics*. London: Routledge and Kegan Paul, 156–198.
- Rosenkrantz, R. (1981), *Foundations and Applications of Inductive Probability*. Atascadero, CA: Ridgeview Press.
- Salmon, W. (1988), "Dynamic Rationality: Propensity, Probability and Credence", in J. H. Fetzer (ed.), *Probability and Causality*. Dordrecht: D. Reidel, 3–40.
- Savage, L. (1971), "Elicitation of Personal Probabilities", *Journal of the American Statistical Association* 66: 783–801.
- Scott, D. (1964), "Measurement Structures and Linear Inequalities", *Journal of Mathematical Psychology* 1: 233–247.
- Seidenfeld, T. (1985), "Calibration, Coherence, and Scoring Rules", *Philosophy of Science* 52: 274–294.
- Shimony, A. (1988), "An Adamite Derivation of the Calculus of Probability", in J. H. Fetzer (ed.), *Probability and Causality*. Dordrecht: D. Reidel, 151–161.
- Skyrms, B. (1984), *Pragmatics and Empiricism*. New Haven: Yale University Press.
- Smith, M. (1987), "The Humean Theory of Motivation", *Mind* 96: 36–61.
- van Fraassen, B. (1983), "Calibration: A Frequency Justification for Personal Probability", in R. Cohen and L. Laudan (eds.), *Physics Philosophy and Psychoanalysis*. Dordrecht: D. Reidel, 295–319.
- Velleman, J. D. (1996), "The Possibility of Practical Reason," *Ethics* 106: 694–726.