The Annals of Statistics
1992, Vol. 20, No. 4, 1737–1767

INDETERMINATE PROBABILITIES ON FINITE SETS[1]

BY ROBERT F. NAU

Duke University

This paper presents a quasi-Bayesian model of subjective uncertainty in which beliefs which are represented by lower and upper probabilities qualified by numerical confidence weights. The representation is derived from a system of axioms of binary preferences which differs from standard axiom systems insofar as completeness is not assumed and transitivity is weakened. Confidence-weighted probabilities may be elicited through the acceptance of bets with limited stakes, a generalization of the operational method of de Finetti. The model is applicable to the reconciliation of inconsistent probability judgments and to the sensitivity analysis of Bayesian decision models.

1. Introduction. The determinacy of personal probabilities has long been controversial among both adherents and critics of the Bayesian viewpoint. Savage (1972), page 58, observes that "there seem to be some probability relations about which we feel relatively 'sure' as compared with others. When our opinions, as reflected in real or envisaged action, are inconsistent, we sacrifice the unsure opinions to the sure ones." Yet his theory provides no measure of the "surety" of a probability nor any prescription for reconciling inconsistency. The intuitive conviction that personal probabilities cannot always be quantified exactly has been cited by Berger (1984), page 64, and Leamer (1986), page 219, as an impediment to the acceptance of Bayesian inference methods, and Berger uses this point to argue for Bayesian methods which are robust against misspecifications of the prior distribution.

The quality of being unsure about a probability cannot itself be satisfactorily modeled by another probability distribution, as the expectation of a second-order distribution is behaviorally indistinguishable from a first-order probability. [Savage attributes this observation to Woodbury; see also Marschak, Borach, Chernoff, DeGroot, Dorfman, Edwards, Ferguson, Miyasawa, Randolph, Savage, Schlaifer and Winkler (1975).] A useful alternative is to represent a subject's beliefs by a convex set of distributions rather than a unique distribution, leading to distinct lower and upper probabilities for events. [An extensive treatment is given by Walley (1991). Special cases are considered by DeRobertis and Hartigan (1981) and Wasserman (1990). There are also varieties of lower and upper probabilities which are not envelopes of convex sets, but these will not be discussed here.] Some authors consider these

Received September 1990; revised August 1991.

[1]This material is based upon work supported by NSF Grant SES-85-11675 and by the Business Associates Fund at the Fuqua School of Business.

AMS 1980 *subject classifications*. Primary 62A15; secondary 90A05, 04A72.

Key words and phrases. Subjective probability, lower and upper probabilities, confidence-weighted probabilities, second-order probabilities, coherence, incompleteness, fuzzy sets.

sets of distributions to represent beliefs which are merely imprecise—that is, theoretically determinate but measured with error. Others consider them to represent beliefs which are intrinsically indeterminate due to limits of rationality or to insufficient quantity or quality of information [Koopman (1940), Smith (1961), Good (1962), Suppes (1974), Kyburg (1974), Williams (1976), Levi (1980), Walley (1982), Seidenfeld, Kadane and Schervish (1989) and Seidenfeld, Schervish and Kadane (1990)]. Competing views are discussed by Levi (1985) and Walley (1991).

As a primitive representation of beliefs, the set-valued probability model is obtained from the standard point-valued model by dropping the axiom of completeness—an axiom many feel to be normatively uncompelling—and in most respects it provides a satisfactory description of indeterminacy. Still, a few intuitive and practical problems remain. To represent a subject's belief in the occurrence of an event by a unique probability interval $[p, q]$ is to endow each of the interval's endpoints with much the same quality of surety as a point probability, whereas intuition suggests that they, too, may be somewhat indeterminate. ["The inequalities themselves have fuzziness," as Good (1962) observes.] The assumption that sets of probabilities are uniquely determined also leaves the following practical problem: If an assessment of lower and upper probabilities turns out to be inconsistent, there is still no basis for "sacrificing the unsure opinions to the sure ones" in order to reconcile the inconsistency.

Parametric families of sets of distributions are widely used in Bayesian sensitivity analysis and robust inference, and various authors have explored the idea that such families might be considered to represent second-order indeterminacy in beliefs. Good (1962) suggests that lower and upper probabilities might be qualified by higher-order distributions—at least conceptually—although Woodbury's objection is seemingly applicable here as well. Gärdenfors and Sahlin (1982, 1983) introduce the notion of an "epistemic reliability" function defined over a set of probability distributions, which yields larger or smaller sets of probabilities as a reliability threshold is varied. Watson, Weiss and Donnell (1979) and Freeling (1980) use "membership" functions to represent probabilities as fuzzy numbers, according to Zadeh's (1965) theory of fuzzy sets. In all of these models an indeterminate probability is described by a general unimodal function on the unit interval rather than a rectangular indicator function. However, the postulated function has no clear-cut behavioral implications, and therefore it is not apparent how to elicit its parameters.

This paper presents a new theory of indeterminate sets of probabilities which addresses the issues raised above. Beliefs are described by lower and upper probabilities qualified by numerical *confidence weights*, which provide a basis for making tradeoffs among probability judgments in order to reconcile inconsistency or perform sensitivity analysis. Confidence weights are operationally defined as the relative stakes attached to different betting rates under a modification of de Finetti's (1937, 1974) elicitation method. As an illustration, consider a subject who is asked to reveal his probability for an event by announcing bets he would be willing to accept at the discretion of an opponent. Suppose that, under the rules of the elicitation game, the subject is not required to determine a rate at which he would indifferently bet on or against the event—thus, he is allowed to announce distinct lower and upper betting rates—and furthermore he is allowed to limit the stakes for which he will bet at a given rate. For example, he might announce more than one lower or upper betting rate for the same event, each with its own limit on stakes. (Presumably he will accept higher stakes for more favorable rates.) The opponent is then allowed to enforce a single bet, or else a convex combination, from among all those offered.

Under these circumstances, suppose the subject offers to accept any one of the following four bets: \$100 *on* the occurrence of an event at a rate of 0.4, \$50 on its occurrence at a rate of 0.5, \$50 *against* its occurrence at a rate of 0.7 or \$100 against its occurrence at a rate of 0.8. [If he bets $\$c$ *on* the event at rate p, the subject loses $\$c$ if the event does not occur and wins $\$c(1-p)/p$ if it does occur. If he bets $\$c$ *against* the event at rate q, the subject loses $\$c$ if the event does occur and wins $\$cq/(1-q)$ if it does not.] The subject thereby reveals that, with some degree of confidence, he places the probability of the event in the interval $[0.5, 0.7]$, and with "twice as much confidence" he places it in the interval $[0.4, 0.8]$: His relative confidence in a betting rate is equated with the relative amount of money he risks losing. Various psychological mechanisms can be proposed to explain this assignment of nontrivial confidence weights to probabilities when the rules of elicitation allow it—for example, it is an appropriate response to an environment in which "people accepting bets against our individual have better information than he has." [See de Finetti (1937), footnote on page 62 of the (1980) edition; also de Finetti (1974), page 93. This interpretation is discussed by Nau (1989, 1990).] However, for the purposes of this paper, judgments of confidence will be considered as behavioral primitives, comparable to numerical judgments of probabilities themselves under de Finetti's original definition.

The properties of confidence-weighted lower and upper probabilities will be deduced from a system of axioms of binary preferences among monetary lotteries over finite state spaces, which differs from standard axiom systems insofar as the assumption of completeness is abandoned and the assumption of transitivity is weakened. [In contrast, Fishburn (1983a, b) models ambiguous beliefs by abandoning transitivity while preserving completeness.] However, de Finetti's axiom of coherence—the avoidance of sure loss—remains as a central requirement of rationality. It will be shown that an assessment of confidence-weighted probabilities for a finite state space is described by a concave function on the probability simplex, and the degree of belief in a single event is described by a unimodal function on the unit interval. These are technically *Bayes risk* functions summarizing the decision problems that the subject's assessment presents to a betting opponent. However, they also have many of the postulated properties of Gärdenfors and Sahlin's epistemic reliability functions and Watson, Weiss and Donnell's and Freeling's membership functions; and they parameterize families of nested, convex sets of distributions which are formally similar to classes of "contaminated" distributions used in robust Bayesian analysis [e.g., Berger (1984) and Berger and Berlinear (1986).]

The novelty of this representation of second-order indeterminacy is that it is linked to material betting behavior and possesses definite rules of inference which follow from the underlying axioms. These, of course, are generalizations of the corresponding rules of inference for point-valued and interval-valued probabilities. On a practical level, it justifies the use of simple linear-programming methods for reconciling incoherence and performing sensitivity analysis in Bayesian decision models [Nau (1989)].

The next section of the paper introduces axioms of confidence-weighted preferences among monetary lotteries, and Section 3 establishes basic duality results analogous to de Finetti's coherence theorem. Section 4 specializes these results for assessments of confidence-weighted probabilities; Section 5 gives an example of inference; and Section 6 discusses potential applications. Proofs are given in Appendix 1 and a computational algorithm is given in Appendix 2.

2. Confidence-weighted preferences among monetary lotteries. Consider a finite set Θ of states of nature with respect to which we wish to characterize an individual's beliefs. Let $\mathbf{X}$, $\mathbf{Y}$, $\mathbf{X}_1$, $\mathbf{Y}_1$, etc., denote bounded random variables over Θ, interpreted as payoff vectors of monetary lotteries, and suppose that an individual ("the subject," "he") reveals his beliefs by asserting preferences among a finite number of pairs of such lotteries. Let $\mathscr{A}_n$ denote the subject's assertion that $\mathbf{X}_n \succsim \mathbf{Y}_n$ ("$\mathbf{X}_n$ is weakly preferred to $\mathbf{Y}_n$") for some lotteries $\mathbf{X}_n$ and $\mathbf{Y}_n$; let $\mathscr{N} = \{1, \dots, N\}$ denote the set of all n for which such direct assertions are made; and let $\mathscr{A}_{\mathscr{N}} = \wedge_{n \in \mathscr{N}} \mathscr{A}_n$, where "$\wedge$" denotes conjunction. $\mathscr{A}_{\mathscr{N}}$ will be called the subject's *assessment* with respect to Θ. We will be interested in judging whether such an assessment is internally consistent and, if so, in determining what other preferences may be inferred from it under suitable axioms of rationality.

Lotteries are represented by vectors in $\mathbb{R}^M$, where $M = |\Theta|$, and will be interpreted as potential changes in the wealth distribution of the subject; the payoff of lottery $\mathbf{X}$ in state $\theta \in \Theta$ will be denoted $X(\theta)$. $\mathbf{0}$ and $\mathbf{1}$ will denote the vectors in $\mathbb{R}^M$ whose elements are identically 0 and 1, respectively, and the symbols $\mathbf{E}$, $\mathbf{F}$, etc., will denote events (subsets of Θ) and also the indicator functions thereof. Thus, $E(\theta) = 1$ if the event $\mathbf{E}$ includes state θ, and $E(\theta) = 0$ otherwise; $\mathbf{1} - \mathbf{E}$ is the complement of $\mathbf{E}$, which will also be denoted as $\overline{\mathbf{E}}$. Addition and multiplication of lotteries by scalars and indicator functions are defined pointwise. For example, $\alpha\mathbf{X} + (1 - \alpha)\mathbf{Y}$ is the lottery paying $\alpha X(\theta) + (1 - \alpha)Y(\theta)$ in state θ, and $\mathbf{EX} + (\mathbf{1} - \mathbf{E})\mathbf{Y}$ is the lottery paying $E(\theta)X(\theta) + (1 - E(\theta))Y(\theta)$ in state θ. The latter is a *conditional lottery* paying $\mathbf{X}$ if $\mathbf{E}$ occurs and $\mathbf{Y}$ otherwise. $[\mathbf{X}]_{\min}$ is defined as the minimum element of $\mathbf{X}$:

$$[\mathbf{X}]_{\min} \equiv \min_{\theta \in \Theta} X(\theta).$$

"$\geq^*$" and "$>^*$" denote nonstrict and strict dominance (vector inequality) between lotteries:

$$\mathbf{X} \geq^* \mathbf{Y} \quad \text{if } [\mathbf{X} - \mathbf{Y}]_{\min} \geq 0$$

and

$$\mathbf{X} >^* \mathbf{Y} \quad \text{if } [\mathbf{X} - \mathbf{Y}]_{\min} > 0.$$

The symbol $\mathbf{B}$ will be used to denote a difference of two lotteries—for example, $\mathbf{B} = \mathbf{X} - \mathbf{Y}$. If $[\mathbf{B}]_{\min} < 0$, then $\mathbf{B}^*$ denotes the normalization of $\mathbf{B}$ so that its minimum element is -1:

$$\mathbf{B}^* \equiv \frac{\mathbf{B}}{-[\mathbf{B}]_{\min}},$$

whence $[\mathbf{B}^*]_{\min} = -1$ by definition. If $\mathscr{Q}$ is any collection of vectors in $\mathbb{R}^M$, CONV$\{\mathscr{Q}\}$ denotes the convex hull of $\mathscr{Q}$, and CONV$^+\{\mathscr{Q}\}$ denotes the sum of CONV$\{\mathscr{Q}\}$ and the nonnegative orthant:

$$\text{CONV}^+\{\mathscr{Q}\} \equiv \{\mathbf{Y} | \mathbf{Y} \geq^* \mathbf{Z} \text{ for some } \mathbf{Z} \in \text{CONV}\{\mathscr{Q}\}\}.$$

The following axioms will be assumed to govern preference:

A.1. (Reflexivity): $\mathbf{X} \succsim \mathbf{X}$.
A.2. (Dominance): $\{\mathbf{X} \succsim \mathbf{Y} \wedge \mathbf{X}' \geq^* \mathbf{X}\} \Rightarrow \mathbf{X}' \succsim \mathbf{Y}$.
A.3. (Cancellation): $\{\mathbf{X} \succsim \mathbf{Y} \wedge \mathbf{X} - \mathbf{Y} = \mathbf{X}' - \mathbf{Y}'\} \Rightarrow \mathbf{X}' \succsim \mathbf{Y}'$.
A.4. (Convexity): $\{\mathbf{X} \succsim \mathbf{Y} \wedge \mathbf{X}' \succsim \mathbf{Y}\} \Rightarrow \alpha\mathbf{X} + (1 - \alpha)\mathbf{X}' \succsim \mathbf{Y} \; \forall \, \alpha \in (0, 1)$.
A.5. (Coherence): $\mathbf{X} >^* \mathbf{Y} \Rightarrow$ Not $\mathbf{Y} \succsim \mathbf{X}$.

These are similar to axioms for preferences among monetary lotteries on finite sets which have been used by de Finetti (1937), Buehler (1976), Williams (1976) and Walley (1982). The dominance axiom implies that more is preferred to less; and cancellation implies that preferences between lotteries do not depend on absolute levels of wealth, only on statewise differences. The latter assumption is reasonable for changes in wealth which are relatively "small." The convexity assumption replaces—and weakens—the usual assumptions of linearity and transitivity. Like Williams and Walley (but unlike de Finetti) we do not assume *completeness* of preferences, that is, we do not assume that, for any $\mathbf{X}$ and $\mathbf{Y}$, either $\mathbf{X} \succsim \mathbf{Y}$ or $\mathbf{Y} \succsim \mathbf{X}$, or both. No Archimedian axiom is included here because only finitely generated preference structures will be considered.

The axioms imply that preferences can be *diluted*: $\mathbf{X} \succsim \mathbf{Y} \Rightarrow \alpha\mathbf{X} \succsim \alpha\mathbf{Y}$ for $\alpha \in (0, 1)$ by successive application of A.1, A.4 and A.3. However, because only convexity rather than linearity is assumed, they cannot necessarily be *undiluted*: The preceding implication does not hold for $\alpha > 1$. Correspondingly, the conventional transitivity property does not hold, but a diluted form of transitivity holds instead. A.3 and A.4 together imply the stronger convexity property $\{\mathbf{X} \succsim \mathbf{Y} \wedge \mathbf{X}' \succsim \mathbf{Y}'\} \Rightarrow \alpha\mathbf{X} + (1 - \alpha)\mathbf{X}' \succsim \alpha\mathbf{Y} + (1 - \alpha)\mathbf{Y}'$ for $\alpha \in (0, 1)$, so that $\{\mathbf{X}_1 \succsim \mathbf{X}_2 \wedge \mathbf{X}_2 \succsim \mathbf{X}_3\} \Rightarrow \frac{1}{2}\mathbf{X}_1 + \frac{1}{2}\mathbf{X}_2 \succsim \frac{1}{2}\mathbf{X}_2 + \frac{1}{2}\mathbf{X}_3$, whence by cancellation we obtain only $\frac{1}{2}\mathbf{X}_1 \succsim \frac{1}{2}\mathbf{X}_3$ instead of the usual $\mathbf{X}_1 \succsim \mathbf{X}_3$. More generally, the same construction yields

$$(2.1)\qquad \{\mathbf{X}_1 \succsim \mathbf{X}_2 \wedge \mathbf{X}_2 \succsim \mathbf{X}_3 \wedge \cdots \wedge \mathbf{X}_k \succsim \mathbf{X}_{k+1}\} \Rightarrow \frac{1}{k}\mathbf{X}_1 \succsim \frac{1}{k}\mathbf{X}_{k+1}.$$

Lest the reader be alarmed at this weakening of transitivity—normally considered a compelling axiom of preference—it should be emphasized that the binary relation axiomatized here is not merely a qualitative ordering relation. The lack of an undilution property means that between any two lotteries $\mathbf{X}$ and $\mathbf{Y}$ there is not merely a direction of preference, but also a secondary, quantitative attribute which later will be termed "confidence," and which is proportional to the maximum value of α for which $\alpha\mathbf{X} \succsim \alpha\mathbf{Y}$. It is this secondary attribute, rather than the qualitative direction of preference, which fails to be conserved in chains of inference such as (2.1). In the presence of the cancellation assumption, we could not assume undiluted transitivity without implying the undilution of preferences, since $\mathbf{X} \succsim \mathbf{Y} \Rightarrow \{\mathbf{X} + \mathbf{X} \succsim \mathbf{X} + \mathbf{Y} \wedge \mathbf{X} + \mathbf{Y} \succsim \mathbf{Y} + \mathbf{Y}\} \Rightarrow 2\mathbf{X} \succsim 2\mathbf{Y}$. The absence of an undilution property (or, equivalently, the dilution of transitivity) distinguishes the model of preferences presented here from Aumann's (1962) utility theory *sans* completeness, Buehler's (1976) model of coherent preferences, and the lower and upper probability models of Smith (1961), Williams (1976), Levi (1980) and Walley (1982, 1991).

If the chain of weak preferences on the left-hand side of (2.1) is converted to a cycle by appending $\mathbf{X}_{k+1} \succsim \mathbf{X}_1$, then a cycle of diluted preferences is implied in the other direction, namely $(1/k)\mathbf{X}_j \succsim (1/k)\mathbf{X}_{j-1}$ for all j. For the reasons just noted, these cannot necessarily be undiluted to infer $\mathbf{X}_j \succsim \mathbf{X}_{j-1}$. However, it is still forbidden to have a cycle of preferences in which at least one preference is strict, under a natural definition of strict preference: $\mathbf{X} \succ \mathbf{Y}$ if $\mathbf{X} \succsim \mathbf{Y}$ and not $\alpha\mathbf{Y} \succsim \alpha\mathbf{X}$ for every $\alpha > 0$. Under this definition, strict preference is dilutable—that is, $\mathbf{X} \succ \mathbf{Y} \Rightarrow \alpha\mathbf{X} \succ \alpha\mathbf{Y}$ for $\alpha \in (0,1)$—and A.5 is equivalent (given the other axioms) to $\mathbf{X} >^* \mathbf{Y} \Rightarrow \mathbf{X} \succ \mathbf{Y}$. If preferences could be undiluted, this definition would reduce to the more familiar one, namely the asymmetric part of $\succsim$ ($\mathbf{X} \succsim \mathbf{Y}$ and not $\mathbf{Y} \succsim \mathbf{X}$). Absent the undilution property, the asymmetric part of $\succsim$ by itself is not a useful definition of strict preference, because it is possible to have $\mathbf{X} \succsim \mathbf{Y}$ and $(1-\varepsilon)\mathbf{Y} \succsim (1-\varepsilon)\mathbf{X}$ but not $\mathbf{Y} \succsim \mathbf{X}$ for arbitrarily small ε. Furthermore, the asymmetric part of $\succsim$ is not dilutable.

Under the axioms given above, the implications of the subject's assessment can be summarized as follows.

DEFINITIONS. A preference relation $\succsim$ on $\mathbb{R}^M$ is a *convex extension* of $\mathscr{A}_{\mathscr{N}}$ if $\mathbf{X}_n \succsim \mathbf{Y}_n$ for all $n \in \mathscr{N}$ and $\succsim$ satisfies A.1–A.4. It is a *coherent extension* if it also satisfies A.5. $\mathscr{A}_{\mathscr{N}}$ is *coherent* if it has a coherent extension.

THEOREM 1. *Let*

$$\mathscr{B}_{\mathscr{N}} \equiv \mathrm{CONV}^+\{\mathbf{0}, \mathbf{B}_1, \ldots, \mathbf{B}_N\},$$

where $\mathbf{B}_n \equiv \mathbf{X}_n - \mathbf{Y}_n \ \forall\ n \in \mathscr{N}$. *Then*:

(i) *The minimal convex extension of* $\mathscr{A}_{\mathscr{N}}$ *is given by* $\mathbf{X} \succsim \mathbf{Y} \Leftrightarrow \mathbf{X} - \mathbf{Y} \in \mathscr{B}_{\mathscr{N}}$.

(ii) $\mathscr{A}_{\mathscr{N}}$ *is coherent if and only if* $\mathscr{B}_{\mathscr{N}}$ *does not contain a strictly negative vector*.

Henceforth, " $\succsim$ " will be used to denote the subject's directly asserted preferences as well as their minimal convex extension. For example, $\mathscr{A}_n = \mathbf{X}_n \succsim \mathbf{Y}_n$ indicates that $\mathscr{A}_n$ is an assertion of preference for $\mathbf{X}_n$ over $\mathbf{Y}_n$, and $\mathscr{A}_{\mathscr{N}} \Rightarrow \mathbf{X} \succsim \mathbf{Y}$ ("$\mathscr{A}_{\mathscr{N}}$ implies that $\mathbf{X}$ is preferred to $\mathbf{Y}$") indicates that the minimal convex extension of $\mathscr{A}_{\mathscr{N}}$ satisfies $\mathbf{X} \succsim \mathbf{Y}$.

Note that the set $\mathscr{B}_{\mathscr{N}}$ consists of (i) the zero vector, (ii) the lottery differences $\mathbf{X}_n - \mathbf{Y}_n$ for all $n \in \mathscr{N}$, (iii) all convex combinations of those preceding and (iv) all vectors that dominate those preceding: It is a convex polyhedron in $\mathbb{R}^M$ having the nonnegative orthant as its cone of recession [Rockafellar (1970)]. $\mathscr{B}_{\mathscr{N}}$ will be called the set of *acceptable gambles*, following de Finetti's operational view that an assertion of preference $\mathbf{X} \succsim \mathbf{Y}$ implies a willingness to exchange $\mathbf{Y}$ for $\mathbf{X}$, which is equivalent to accepting a gamble whose payoff vector is $\mathbf{X} - \mathbf{Y}$, at the discretion of an opponent. The decision problem for the opponent ("she") is to select a gamble from this set to enforce. By comparison, the set of acceptable gambles is a convex cone under Smith's operational definition of lower and upper probabilities, and it is a half-space

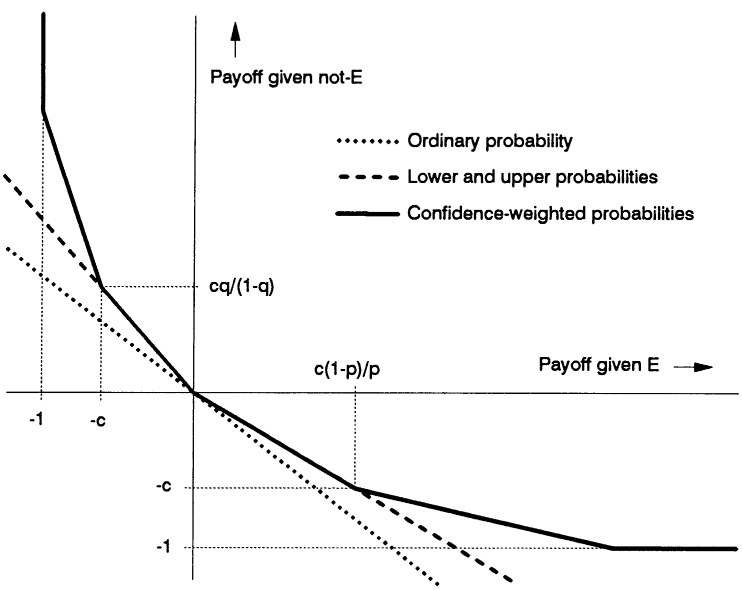

FIG. 1. *Frontiers of sets of acceptable gambles corresponding to assessments of an ordinary probability, a lower and upper probability and confidence-weighted probabilities. p and q are confidence-weighted lower and upper probabilities with confidence c.*

under de Finetti's definition of ordinary probabilities. (Under de Finetti's and Smith's definitions, acceptable gambles are additive, not merely convexifiable.) These differences are illustrated in Figure 1 for the case $M = 2$, where $\Theta = \{\mathbf{E}, \overline{\mathbf{E}}\}$ for some event $\mathbf{E}$.

The secondary attribute of preference noted above will now be made explicit. Let $\mathbf{X}$ and $\mathbf{Y}$ be lotteries for which $[\mathbf{X} - \mathbf{Y}]_{\min} < 0$, and suppose the subject asserts that $\alpha\mathbf{X}$ is preferred to $\alpha\mathbf{Y}$ for some $\alpha > 0$, meaning that he will accept a gamble whose payoff vector is $\alpha(\mathbf{X} - \mathbf{Y})$. Then $c = -\alpha[\mathbf{X} - \mathbf{Y}]_{\min}$ is the maximum amount of money the subject might lose if his acceptance of this gamble were enforced by the opponent. Using the normalization operation introduced above, this gamble can be equivalently expressed as $c(\mathbf{X} - \mathbf{Y})^*$. Thus, c is the size of the *stake* the subject is willing to risk in accepting a gamble proportional to $(\mathbf{X} - \mathbf{Y})$, and it will be defined as the amount of *confidence* with which he prefers $\mathbf{X}$ to $\mathbf{Y}$:

DEFINITION. $\mathbf{X} \succsim_c \mathbf{Y}$ ("$\mathbf{X}$ is preferred to $\mathbf{Y}$ with confidence c") if $\alpha\mathbf{X} \succsim \alpha\mathbf{Y}$ and $c = -\alpha[\mathbf{X} - \mathbf{Y}]_{\min}$.

From Theorem 1, it follows that $\mathscr{A}_{\mathscr{N}} \Rightarrow \mathbf{X} \succsim_c \mathbf{Y}$ is and only if $c(\mathbf{X} - \mathbf{Y})^* \in \mathscr{B}_{\mathscr{N}}$. The subject's assessment (as well as its inferences) will hereafter be recoded entirely in terms of confidence-weighted preferences. For example, we will write $\mathscr{A}_n = \mathbf{X}_n \succsim_{c_n} \mathbf{Y}_n$ (meaning that $\alpha_n \mathbf{X}_n$ is asserted to be preferred to $\alpha_n \mathbf{Y}_n$, and $c_n = -\alpha_n[\mathbf{X}_n - \mathbf{Y}_n]_{\min}$) for every n. $\succsim_c$ will be considered as the minimal convex extension of the subject's assessment of confidence-weighted preferences. In the terminology of Luce and Narens (1985), $\succsim_c$ is a *relational structure* comprising an uncountable infinity of binary relations indexed by c. Note that $\succsim_c$, unlike $\succsim$, possesses the undilution property:

$$\mathbf{X} \succsim_c \mathbf{Y} \Leftrightarrow \alpha\mathbf{X} \succsim_c \alpha\mathbf{Y} \qquad \forall\, \alpha > 0.$$

Given a pair of lotteries $\mathbf{X}$ and $\mathbf{Y}$, neither of which dominates the other, we will generally be interested in determining the *maximum* confidence with which it can be inferred that $\mathbf{X}$ is preferred to $\mathbf{Y}$. Note that if preferences could be undiluted, the acceptability of a gamble proportional to $\mathbf{X} - \mathbf{Y}$ would be independent of the stake: The subject would have infinite confidence in his preference for $\mathbf{X}$ over $\mathbf{Y}$ if he had any preference at all. Absent the undilution property, the subject's confidence in his preference for $\mathbf{X}$ over $\mathbf{Y}$ will be finite if $\mathbf{X}$ does not dominate $\mathbf{Y}$. In particular, if we define

$$c_{\mathscr{N}} \equiv \max_{n \in \mathscr{N}} c_n,$$

where c_n is the degree of confidence in assertion $\mathscr{A}_n$, then $c_{\mathscr{N}}$ is an upper bound on the confidence of any preferences which may be inferred from the assessment, because inferences are obtained through the formation of convex combinations. For convenience, we will hereafter take $c_{\mathscr{N}}$ as the unit of confidence (i.e., we set $c_{\mathscr{N}} = 1$) so that all confidence weights will lie between 0 and 1.

There are several reasons for identifying the subject's "confidence" in a gamble with the amount of money he is willing to lose, rather than some other measure such as the amount he might win, the length of the payoff vector in Euclidean space or another attribute. First, it will be demonstrated that this definition has mathematically interesting properties which are in accord with intuitive notions of a subject's degree of conviction in his beliefs. Second, it is psychologically convenient: Our goal is to associate a confidence weight with an odds ratio or probability, and the definition of confidence in this way coincides with the familiar decomposition of gambles into "odds" and "stakes." Third, it is closely related to the formal definition of *regret*—that is, the difference between what is actually received and what might have been received had another decision been made—which a number of researchers have identified as a factor in behavior which violates subjective expected utility theory, particularly in situations where the estimation of a probability is subject to second-guessing [Bell (1982) and Loomes and Sugden (1982)]. The subject's confidence in accepting a gamble, as defined here, is the relative amount of regret to which he thereby exposes himself.

Although the focus here is on the modeling of probabilities and not utilities, the axioms given above do not rule out the possibility that the subject is a subjective-expected-utility maximizer with constant absolute risk aversion—that is, exponential utility for money. If this were the case, his apparent confidence in any gamble having both positive and negative payoffs would always be finite merely because of risk aversion, even if his beliefs were highly determinate. [Leamer (1986) discusses the phenomenon of lower and upper probabilities in precisely these terms.] Specifically, we would expect his confidence in any reasonable gamble to be of the same order of magnitude as his *risk tolerance*, and we would also expect to see certain patterns among the gambles he would accept on different events. [If the subject's utility-for-money function has the exponential form $u(x) = 1 - \exp(-x/T)$, the constant T is his risk tolerance. Roughly speaking, it is the amount of money such that he would be indifferent to accepting a gamble offering equal chances of winning that amount or losing half that amount.] For example, he would be able to weakly order events according to likelihood, and we would expect him to accept the gambles $\mathbf{E} - p\mathbf{1}$ and $\mathbf{E}' - p\mathbf{1}$ with equal confidence for every value of p if the events $\mathbf{E}$ and $\mathbf{E}'$ were judged equally likely and had no a priori bearing on his wealth. The viewpoint adopted here, nonetheless, is that the subject's attachment of finite confidence weights to preferences is *not* primarily due to risk aversion, that the amount of money at risk is small in comparison with his risk tolerance (if the latter is finite) and that its value is in fact somewhat arbitrary. Thus, it is assumed that the information conveyed by the confidence weights inheres in the ratios $c_n/c_{\mathscr{N}}$ rather than the magnitude of $c_{\mathscr{N}}$.

3. Duality results: coherence and inference. The convex set $\mathscr{B}_{\mathscr{N}}$ of acceptable gambles generated by $\mathscr{A}_{\mathscr{N}}$ constitutes the primal geometric representation of the subject's beliefs, and there is a corresponding dual or conjugate representation in the form of a concave function on the simplex of probability distributions on Θ [Rockafellar (1970), Corollary 16.5.1]. In the fundamental coherence theorem for confidence-weighted probabilities to be given below, this function plays the role that is occupied by a unique probability distribution in de Finetti's theory, or by a convex set of "medial" distributions in Smith's theory. This function is in fact the *Bayes risk* function summarizing the statistical decision problem that the subject's assessment presents to his betting opponent. The Bayes risk for a statistical decision problem is conventionally defined as the minimum achievable expected loss for the decision maker (here, the opponent) as a function of her probability distribution, where "loss" is measured in such a way as to be intrinsically nonnegative [DeGroot (1970)].

Let Π denote the simplex of probability distributions on Θ, and let $\pi \mapsto P_\pi(\cdot)$ denote the probability (expectation) assigned to an event (lottery) by the distribution $\pi \in \Pi$. That is,

$$P_\pi(\mathbf{E}) \equiv \sum_{\theta \in \Theta} \pi(\theta) E(\theta)$$

and

$$P_\pi(\mathbf{E}|\mathbf{F}) \equiv P_\pi(\mathbf{EF})/P_\pi(\mathbf{F}) \quad \text{if } P_\pi(\mathbf{F}) > 0.$$

The subject's assertion that $\mathbf{X}_n$ is preferred to $\mathbf{Y}_n$ with confidence c_n means he will accept a gamble whose gain for him is $c_n \mathbf{B}_n^*$, where $\mathbf{B}_n = \mathbf{X}_n - \mathbf{Y}_n$, and whose gain for the opponent is therefore $-c_n \mathbf{B}_n^*$. If the opponent's distribution on Θ is π, she should enforce this gamble only if $P_\pi(\mathbf{B}_n^*) \leq 0$, obtaining an expected gain of $-c_n P_\pi(\mathbf{B}_n^*)$; otherwise, she should decline the gamble, obtaining an expected gain of 0. (Introduction of "the opponent's distribution π" should not be interpreted to mean that we impute to the opponent the very determinacy of beliefs which we deny to the subject; it is merely a technical device whose usefulness will ultimately be justified by the role it plays in representation theorems given below.) Consequently, the maximum achievable expected gain presented to her by $\mathscr{A}_n$ alone is $-c_n \min\{0, P_\pi(\mathbf{B}_n^*)\}$. The opponent's maximum possible gain over all acceptable gambles is the subject's maximum possible loss, which is equal to 1 by our choice of units. Therefore, we define a nonnegative "loss" for the opponent as 1 minus her actual gain, and the "Bayes risk against π induced by $\mathscr{A}_n$," denoted as $\pi \mapsto R_\pi(\mathscr{A}_n)$, is then given by

$$R_\pi(\mathscr{A}_n) = 1 + c_n \min\{0, P_\pi(\mathbf{B}_n^*)\}.$$

The decision problem posed by the entire assessment is a mixture (convex combination) of the problems posed by the separate assertions. In general, the Bayes risk induced by a mixture of decision problems is the pointwise minimum of their separate Bayes risks, since for any π it is always an optimal strategy to take the *pure* decision yielding the smallest expected loss. Therefore, the Bayes risk against π induced by the entire assessment, denoted as

$\pi \mapsto R_\pi(\mathscr{A}_\mathscr{N})$, is given by

$$R_\pi(\mathscr{A}_\mathscr{N}) = \min_{n \in \mathscr{N}} R_\pi(\mathscr{A}_n) = \min_{n \in \mathscr{N}} 1 + c_n \min\{0, P_\pi(\mathbf{B}_n^*)\}.$$

Being the pointwise minimum of a finite collection of affine functions, it is piecewise linear and concave on the simplex.

Let $\mathscr{A}$ denote a hypothetical additional assertion that $\mathbf{X}$ is preferred to $\mathbf{Y}$ with confidence c, for given values of $\mathbf{X}$, $\mathbf{Y}$ and c; and let $\mathbf{B} = c(\mathbf{X} - \mathbf{Y})^*$. In other words, $\mathbf{B}$ is the gamble whose acceptance would operationally define the assertion $\mathscr{A}$. By the same reasoning as above, the Bayes risk which would be induced by the assertion $\mathscr{A}$ alone is given by

$$R_\pi(\mathscr{A}) = 1 + c \min\{0, P_\pi(\mathbf{B}^*)\}.$$

However, recall from Theorem 1 that $\mathscr{A}$ is already implied by $\mathscr{A}_\mathscr{N}$ if and only if $\mathbf{B} \in \mathscr{B}_\mathscr{N}$. In these terms, the "fundamental theorem" of confidence-weighted preferences is as follows.

THEOREM 2.

(i) $\mathscr{A}_\mathscr{N}$ *is coherent if and only if* $R_\pi(\mathscr{A}_\mathscr{N}) = 1$ *for some* $\pi \in \Pi$.
(ii) $\mathscr{A}_\mathscr{N} \Rightarrow \mathscr{A}$ *if and only if* $R_\pi(\mathscr{A}_\mathscr{N}) \le R_\pi \mathscr{A}) \;\forall\; \pi \in \Pi$.

Note that $R_\pi(\mathscr{A}_\mathscr{N}) = 1$ implies $P_\pi(\mathbf{B}_n) \ge 0$, or equivalently $P_\pi(\mathbf{X}_n) \ge P_\pi(\mathbf{Y}_n)$, for all $n \in \mathscr{N}$. Hence, coherence requires the existence of a probability distribution under which every asserted preference is in the direction of increasing expected value.

The construction of the function $R_\pi(\mathscr{A}_\mathscr{N})$ can be visualized as follows. Imagine an M-dimensional frosted layer cake whose base is the simplex Π (which has dimension $M - 1$) and whose height (in the z direction) is initially equal to 1 everywhere. Then, for each preference assertion of the form $\mathbf{X}_n \succsim_{c_n} \mathbf{Y}_n$, imagine making a downward, angled cut into the cake and removing the material above the cut. In particular, let the nth cutting plane be the set of (π, z) satisfying $z = 1 + c_n P_\pi(\mathbf{B}_n^*)$, where $\mathbf{B}_n = \mathbf{X}_n - \mathbf{Y}_n$. For $c_n \in [0, 1]$, such a cut does not intersect the base of the cake: It enters through the top along the line where $P_\pi(\mathbf{X}_n) = P_\pi(\mathbf{Y}_n)$ and exits through the side along which $P_\pi(\mathbf{X}_n) < P_\pi(\mathbf{Y}_n)$ at a height of $1 - c_n$ above the base. The graph of $R_\pi(\mathscr{A}_\mathscr{N})$ on Π is the upper surface of the whittled-down cake that remains after all N cuts have been made. The assessment is coherent if any of the frosted top surface remains—that is, if the cake still attains an altitude of 1 at some point. The assertion $\mathbf{X} \succsim_c \mathbf{Y}$ can be inferred from the assessment if its own cutting plane lies above or tangent to the remainder of the cake—that is, if it would not remove any additional material.

4. Confidence-weighted probabilities. The case in which an assessment refers directly to probabilities is of special interest. So, for events $\mathbf{E}$ and $\mathbf{F}$ and numbers p and c between 0 and 1, we introduce the following:

DEFINITIONS. The conditional probability of $\mathbf{E}$ given $\mathbf{F}$ is at least p with confidence c, or equivalently, (p, c) is a *confidence-weighted lower probability* for $\mathbf{E}|\mathbf{F}$, if $\mathbf{EF} \succsim_c p\mathbf{F}$. Similarly, (q, c) is a *confidence-weighted upper probability* for $\mathbf{E}|\mathbf{F}$ if $\mathbf{EF} \precsim_c q\mathbf{F}$, or equivalently, $(\mathbf{1} - \mathbf{E})\mathbf{F} \succsim_c (1 - q)\mathbf{F}$.

The assertion of (p, c) as a confidence-weighted lower probability for $\mathbf{E}|\mathbf{F}$ implies acceptance of the gamble whose payoff vector is $c(\mathbf{EF} - p\mathbf{F})^*$. Since $\mathbf{EF} - p\mathbf{F}$ is normalized upon division by p, this gamble can be rewritten as $(c/p)(\mathbf{E} - p\mathbf{1})\mathbf{F}$. Confidence-weighted probabilities defined in this way are a generalization of lower and upper probabilities as operationally defined by Smith (1961), which in turn are generalizations of ordinary probabilities as operationally defined by de Finetti (1937, 1974). This is illustrated by Figure 1, in which p and q are confidence-weighted lower and upper probabilities with confidence c for the unconditional event $\mathbf{E}$. The properties of confidence-weighted probabilities are implicit in the results of Theorem 2, but this section will recast those results in terms that highlight their relationship with other representations of beliefs.

When beliefs are represented by ordinary probabilities (or intervals of probabilities), an assessment is usually made by eliciting probability judgments with respect to some finite set of "source" events. These judgments determine a distribution (or convex set of distributions) in the simplex, which is subsequently marginalized to yield inferences concerning the probabilities of other, "target" events. A similar procedure exists for confidence-weighted probabilities: In this case, it is the Bayes risk function on Π that is marginalized. For example, suppose (p, c) is asserted to be a confidence-weighted (unconditional) lower probability for $\mathbf{E}$, and let x denote a possible value for the *opponent's* marginal probability of $\mathbf{E}$. [That is, $x = P_\pi(\mathbf{E})$, where π denotes the opponent's distribution on Θ.] The associated gamble yields a gain of $-(c/p)(\mathbf{E} - p\mathbf{1})$ to the opponent, and its expected value for her is $-(c/p)(x - p)$, which is positive if $x < p$. Hence, the Bayes risk (which is 1 minus the opponent's maximum achievable expected gain) depends on π only through the marginal probability x. This will be called the "marginal" Bayes risk against x induced by (p, c); it will be denoted as $x \mapsto \mu_x(p, c)$, and it is given by

$$(4.1)\quad \mu_x(p,c) = \begin{cases} 1, & \text{if } x \geq p, \\ 1 + (c/p)(x - p), & \text{otherwise,} \end{cases} = 1 + c\min\left\{0, \frac{x}{p} - 1\right\}.$$

By complementarity, the marginal Bayes risk against x induced by the assertion of (q, c) as a confidence-weighted *upper* probability for $\mathbf{E}$ can be expressed in terms of μ as

$$\mu_{1-x}(1 - q, c) = 1 + c\min\left\{0, \frac{1 - x}{1 - q} - 1\right\}.$$

The function μ will serve as the building block for subsequent constructions: For fixed p and c, the graph of $\mu_x(p, c)$ versus x consists of line segments

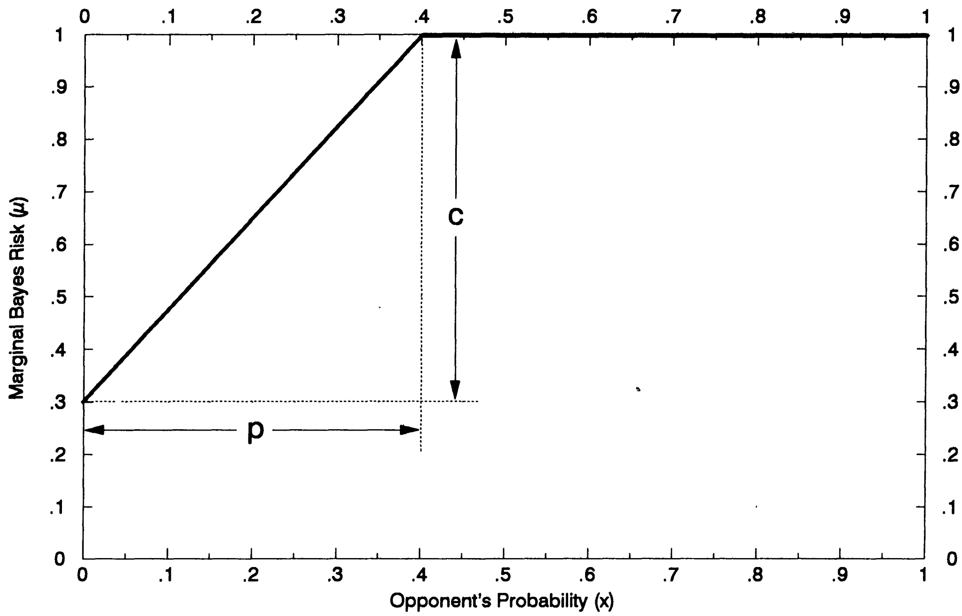

FIG. 2. *Generic marginal Bayes risk* (MBR) *function induced by the assessment of p as a lower probability with confidence c* ($p = 0.4$, $c = 0.7$).

connecting the three points whose (x, y) coordinates are $(0, 1 - c)$, $(p, 1)$ and $(1, 1)$, as shown in Figure 2.

With lower and upper probabilities qualified by confidence weights, the subject may express his uncertainty about an event using several lower and/or upper probabilities at different levels of confidence. For example, he might assert that (p_1, c_1) and (p_2, c_2) are confidence-weighted lower probabilities for the same event, with $c_1 > c_2$ if $p_1 < p_2$. The lesser of two lower probabilities for the same event must have the greater confidence in order to be informative; otherwise, its implied gamble would be dominated for the opponent by the other gamble. Similarly, the greater of two values for an upper probability must have the greater confidence. Thus, the subject's uncertainty may be represented by a sequence of nested intervals $[p_1, q_1] \supset [p_2, q_2] \supset \cdots$, indexed by confidence weights $c_1 > c_2 > \cdots$, rather than by a single point or interval in $[0, 1]$. This nested sequence of intervals may be thought of as a single "fuzzy-edged" probability interval, for reasons to be elaborated below.

First, consider the properties of an assessment given entirely in terms of *unconditional* confidence-weighted probabilities for a set of source events, $\{\mathbf{E}_n | n \in \mathscr{N}\}$. To explicitly accommodate multiple lower and/or upper probabilities for each event, a more detailed notation will be introduced. Let the subject's assessment consist of a conjunction of assertions $\{\mathscr{A}_{nj} | n \in \mathscr{N},\ j \in \mathscr{J}_n\}$, where

$$\mathscr{A}_{nj} = \left\{\mathbf{E}_n \succsim_{c_{nj}} p_{nj}\mathbf{1} \wedge \mathbf{E}_n \precsim_{c_{nj}} q_{nj}\mathbf{1}\right\}.$$

In other words, $\mathscr{A}_{nj}$ is the joint assertion that p_{nj} and q_{nj} are confidence-weighted lower and upper probabilities, respectively, with confidence c_{nj} for the event $\mathbf{E}_n$. (It is not necessary to assume lower and upper probabilities are always assessed in matched pairs at the same level of confidence, but there is no loss of generality in doing so, and it will simplify subsequent comparisons with other models.) Henceforth, let $\mathscr{A}_n = \wedge_{j \in \mathscr{J}_n} \mathscr{A}_{nj}$ and $\mathscr{A}_{\mathscr{N}} = \wedge_{n \in \mathscr{N}} \mathscr{A}_n$. That is, $\mathscr{A}_n$ is the partial assessment referring only to $\mathbf{E}_n$, and $\mathscr{A}_{\mathscr{N}}$ is once again the entire assessment. Given the assessment $\mathscr{A}_{\mathscr{N}}$, we wish to determine the confidence-weighted probabilities which can be inferred for an arbitrary target event $\mathbf{E}$.

The function μ introduced above may be used to construct Bayes risk functions summarizing different levels of the assessment. First, let $x \mapsto r_x(\mathbf{E}_n; \mathscr{A}_{nj})$ denote the "marginal Bayes risk against x as a value for the probability of $\mathbf{E}_n$ that is induced by $\mathscr{A}_{nj}$." Since $\mathscr{A}_{nj}$ is actually the conjunction of two elementary assertions—namely that (p_{nj}, c_{nj}) and (q_{nj}, c_{nj}) are confidence-weighted lower and upper probabilities, respectively, for $\mathbf{E}_n$—the Bayes risk it induces is the pointwise minimum of their separate contributions:

$$\begin{aligned} r_x(\mathbf{E}_n; \mathscr{A}_{nj}) &= \min\{\mu_x(p_{nj}, c_{nj}), \mu_{1-x}(1 - q_{nj}, c_{nj})\} \\ &= 1 + c_{nj} \min\left\{0, \frac{x}{p_{nj}} - 1, \frac{1-x}{1-q_{nj}} - 1\right\}. \end{aligned} \tag{4.2a}$$

The graph of this function on $[0, 1]$ consists of line segments connecting the four points $(0, 1 - c_{nj})$, $(p_{nj}, 1)$, $(q_{nj}, 1)$ and $(1, 1 - c_{nj})$. The pointwise minimum of these functions for all $j \in \mathscr{J}_n$ yields the marginal Bayes risk induced by the partial assessment $\mathscr{A}_n$:

$$r_x(\mathbf{E}_n; \mathscr{A}_n) = \min_{j \in \mathscr{J}_n} r_x(\mathbf{E}_n; \mathscr{A}_{nj}), \tag{4.2b}$$

which is piecewise linear and concave on $[0, 1]$. Next, the Bayes risk induced by $\mathscr{A}_n$ against the full distribution π is obtained by extending $r(\mathbf{E}_n; \mathscr{A}_n)$ to the simplex Π via the mapping $x \mapsto \{\pi | P_\pi(\mathbf{E}_n) = x\}$:

$$\begin{aligned} R_\pi(\mathscr{A}_n) &= r_{x_n(\pi)}(\mathbf{E}_n; \mathscr{A}_n) \\ &= \min_{j \in \mathscr{J}_n} 1 + c_{nj} \min\left\{0, \frac{x_n(\pi)}{p_{nj}} - 1, \frac{1 - x_n(\pi)}{1 - q_{nj}} - 1\right\}, \end{aligned} \tag{4.3}$$

where $x_n(\pi) \equiv P_\pi(\mathbf{E}_n)$. This function is piecewise linear and concave on Π. The Bayes risk against π induced by the entire assessment $\mathscr{A}_{\mathscr{N}}$ is the pointwise minimum of these functions over n:

$$R_\pi(\mathscr{A}_{\mathscr{N}}) = \min_{n \in \mathscr{N}} R_\pi(\mathscr{A}_n). \tag{4.4}$$

Now suppose that for some target event $\mathbf{E}$ it is desired to determine the marginal Bayes risk against x as a value for the probability of $\mathbf{E}$ which is induced by the assessment $\mathscr{A}_{\mathscr{N}}$. Consistent with the notation introduced above, this will be denoted as $x \mapsto r_x(\mathbf{E}; \mathscr{A}_{\mathscr{N}})$. Given knowledge of x (but not necessarily π), it can only be said that the opponent's minimum achievable expected loss is less than or equal to the maximum of the Bayes risk $R_\pi(\mathscr{A}_{\mathscr{N}})$ over all π satisfying $P_\pi(\mathbf{E}) = x$ whence

$$(4.5) \qquad r_x(\mathbf{E}; \mathscr{A}_{\mathscr{N}}) = \max_{x:\, P_\pi(\mathbf{E}) = x} R_\pi(\mathscr{A}_{\mathscr{N}}).$$

Note that if $\mathbf{E} = \mathbf{E}_n$ for some $n \in \mathscr{N}$—that is, if the target event is taken to be one of the source events—it follows by substitution of (4.3) and (4.4) in (4.5) that $r_x(\mathbf{E}; \mathscr{A}_{\mathscr{N}}) \le r_x(\mathbf{E}_n; \mathscr{A}_n)$ for all $x \in [0, 1]$.

When the source and/or target events are *conditional*, the development is only slightly more complex. Let $\mathscr{A}_{nj}$ now denote the joint assertion of (p_{nj}, c_{nj}) and (q_{nj}, c_{nj}) as confidence-weighted lower and upper probabilities for $\mathbf{E}_n$ conditional on $\mathbf{F}_n$. That is,

$$\mathscr{A}_{nj} \equiv \left\{\mathbf{E}_n \mathbf{F}_n \succsim_{c_{nj}} p_{nj} \mathbf{F}_n \wedge \mathbf{E}_n \mathbf{F}_n \precsim_{c_{nj}} q_{nj} \mathbf{F}_n\right\}.$$

The marginal Bayes risk $x \mapsto r_x(\mathbf{E}_n | \mathbf{F}_n; \mathscr{A}_n)$ induced by $\mathscr{A}_n$ depends on the parameters p_{nj}, q_{nj} and c_{nj} in the same way as before:

$$r_x(\mathbf{E}_n | \mathbf{F}_n; \mathscr{A}_n) = \min_{j \in \mathscr{J}_n} \min\left\{\mu_x(p_{nj}, c_{nj}), \mu_{1-x}(1 - q_{nj}, c_{nj})\right\}.$$

Strictly speaking, this should be called the *conditional* Bayes risk induced by $\mathscr{A}_n$ against x as a value for the probability of $\mathbf{E}_n$ given $\mathbf{F}_n$, although we will still refer to it generically as a "marginal" Bayes risk (MBR) function. Now, conditional on the occurrence of $\mathbf{F}_n$, which has probability $P_\pi(\mathbf{F}_n)$, the opponent's minimum achievable expected loss with respect to $\mathscr{A}_n$ alone is the marginal Bayes risk given above; and conditional on $\overline{\mathbf{F}}_n$, which has probability $1 - P_\pi(\mathbf{F}_n)$, her minimum achievable expected loss is 1. (If $\mathbf{F}_n$ fails to occur, the bets on $\mathbf{E}_n$ are called off: This yields the opponent a gain of 0, which is considered a "loss" of 1 relative to her maximum possible gain.) The unconditional Bayes risk against π induced by $\mathscr{A}_n$ is therefore given by the modified extension formula:

$$(4.3') \qquad \begin{aligned} R_\pi(\mathscr{A}_n) &= P_\pi(\mathbf{F}_n) r_{x_n(\pi)}(\mathbf{E}_n | \mathbf{F}_n; \mathscr{A}_n) + (1 - P_\pi(\mathbf{F}_n)) \\ &= 1 - P_\pi(\mathbf{F}_n)\left[1 - r_{x_n(\pi)}(\mathbf{E}_n | \mathbf{F}_n; \mathscr{A}_n)\right] \\ &= \min_{j \in \mathscr{J}_n} 1 + P_\pi(\mathbf{F}_n) c_{nj} \min\left\{0, \frac{x_n(\pi)}{p_{nj}} - 1, \frac{1 - x_n(\pi)}{1 - q_{nj}} - 1\right\}, \end{aligned}$$

where $x_n(\pi) \equiv P_\pi(\mathbf{E}_n | \mathbf{F}_n)$. By inverting this operation, we obtain the corresponding projection formula:

$$(4.5') \qquad r_x(\mathbf{E} | \mathbf{F}; \mathscr{A}_{\mathscr{N}}) = \sup_{\pi:\, P_\pi(\mathbf{E}|\mathbf{F}) = x,\, P_\pi(\mathbf{F}) > 0} 1 - \frac{1 - R_\pi(\mathscr{A}_{\mathscr{N}})}{P_\pi(\mathbf{F})}.$$

The "laws" of confidence-weighted probabilities are now formalized in the following theorem.

THEOREM 3. *For any assessment* $\mathcal{A}_{\mathcal{N}}$ *and any target events* **E** *and* **F**:

(i) $r(\mathbf{E}|\mathbf{F};\mathcal{A}_{\mathcal{N}})$ *is nonnegative, piecewise linear and concave on* $[0,1]$; *furthermore, if* $R_\pi(\mathcal{A}_{\mathcal{N}}) = 1$ *for some* π *such that* $P_\pi(\mathbf{F}) > 0$, *then*

$$\max_{x\in[0,1]} r_x(\mathbf{E}|\mathbf{F};\mathcal{A}_{\mathcal{N}}) = 1.$$

(ii) *Consequently, there exist finite sets of confidence-weighted lower probabilities* $\{(\hat{p}_s, \hat{c}_s),\ s \in S\}$ *and confidence-weighted upper probabilities* $\{(\hat{q}_t, \hat{c}_t),\ t \in T\}$ *such that*

$$r_x(\mathbf{E}|\mathbf{F};\mathcal{A}_{\mathcal{N}}) = \min\left\{\left(\min_{s\in S} \mu_x(\hat{p}_s, \hat{c}_s)\right), \left(\min_{t\in T} \mu_{1-x}(1-\hat{q}_t, \hat{c}_t)\right)\right\}.$$

(iii) $\mathcal{A}_{\mathcal{N}} \Rightarrow \mathbf{EF} \succsim_c p\mathbf{F}$ *if and only if* $r_x(\mathbf{E}|\mathbf{F};\mathcal{A}_{\mathcal{N}}) \le \mu_x(p,c)$ *for all* $x \in [0,1]$.

Hence, the inferences which can be derived from a finite assessment with respect to an arbitrary target event can be summarized by finitely many confidence-weighted probabilities. Note that $\hat{p}$ is the *greatest lower probability with confidence* $\hat{c}$ which can be inferred for $\mathbf{E}|\mathbf{F}$ if

$$\hat{p} = \max\{p | r_x(\mathbf{E}|\mathbf{F};\mathcal{A}_{\mathcal{N}}) \le \mu_x(p,\hat{c})\ \forall\, x \in [0,1]\}. \tag{4.6}$$

In this case, we will say that $(\hat{p}, \hat{c})$ is an *informative* confidence-weighted lower probability for $\mathbf{E}|\mathbf{F}$ in the context of the assessment $\mathcal{A}_{\mathcal{N}}$. Geometrically, $(\hat{p}, \hat{c})$ is informative if the line passing through the points $(0, 1-\hat{c})$ and $(\hat{p}, 1)$ is tangent to the graph of the marginal Bayes risk function $r(\mathbf{E}|\mathbf{F};\mathcal{A}_{\mathcal{N}})$.

More generally, it is possible to elicit and infer lower and upper expectations (previsions) for lotteries with arbitrary payoffs. In this way, families of arbitrarily shaped convex sets of distributions can be determined, overcoming the limitations of "intervalism" discussed by Levi (1980). Let (p, c) be defined as a *confidence-weighted lower conditional expectation* for a lottery **X** given the occurrence of an event **F** if $\mathbf{XF} \succsim_c p\mathbf{F}$. From the undilution property of confidence-weighted preferences, it follows that for any constants α and β with $\beta > 0$, (p, c) is a confidence-weighted lower expectation for $\mathbf{X}|\mathbf{F}$ if and only if $(\alpha + \beta p, c)$ is a confidence-weighted lower expectation for $(\alpha\mathbf{1} + \beta\mathbf{X})|\mathbf{F}$. It therefore suffices to restrict attention to lotteries which have been normalized by positive linear transformations so that their minimum and maximum elements are 0 and 1, respectively. The results of Theorem 3 apply to confidence-weighted expectations if **E** is taken to be a normalized lottery rather than the indicator of an event.

Throughout the sequence of inferential calculations above, the functions $\mu(\cdot)$, $r(\cdot)$ and $R(\cdot)$ play the role of generalized indicator functions on $[0,1]$ or Π, in the sense that the max and min operations performed on them are identical to those which would be performed on indicator functions for convex sets of probabilities induced by an assessment of lower and upper probabilities under the Koopman–Smith–Good model. Consequently, for any target events, the greatest lower and least upper probabilities with nonzero confidence (which are the endpoints of intervals on which the marginal Bayes risk is equal to 1) obey the Koopman–Smith–Good laws of lower and upper probabilities. Since the idea of a generalized indicator function taking on values intermediate between 0 and 1 and obeying max/min rules for union and intersection was the motivation behind Zadeh's (1965) original definition of a fuzzy set, it is suggestive to think of the Bayes risk functions as membership functions for fuzzy convex sets of probabilities. Thus, the subject's assertion of p as a lower probability with confidence c induces a fuzzy subinterval of $[0,1]$ having the "atomic" membership function $\mu(p,c)$ given by (4.1); the intersection of these over all $\{(p_{nj}, c_{nj}) | j \in \mathscr{J}_n\}$ yields a fuzzy subinterval with membership function $r(\mathbf{E}_n; \mathscr{A}_n)$ given by (4.2a) and (4.2b), summarizing all the direct assertions about $\mathbf{E}_n$; extension of this to the simplex via the mapping $x \mapsto \{\pi\colon P_\pi(\mathbf{E}_n | \mathbf{F}_n) = x\}$ yields a fuzzy subset of Π with membership function $R(\mathscr{A}_n)$ given by (4.3); the intersection of these over all n yields a fuzzy convex subset of Π with membership function $R_\pi(\mathscr{A}_{\mathscr{N}})$ given by (4.4), summarizing the entire assessment. Finally, projection of the latter back onto the unit interval via the mapping $\pi \mapsto P_\pi(\mathbf{E})$ yields a fuzzy probability interval for the target event $\mathbf{E}$ whose membership function is $r(\mathbf{E}; \mathscr{A}_{\mathscr{N}})$ given by (4.5), from which informative confidence-weighted probabilities may be extracted via (4.6).

This is literally a "fuzzification" of the Koopman–Smith–Good model; it should not be confused with attempts to link fuzzy set theory to the Dempster–Shafer theory of belief functions [DuBois and Prade (1989)] or to nonadditive probability theory based on Choquet integration [Wakker (1990)]. Rather, it is formally similar to Watson, Weiss and Donnell's (1979) and Freeling's (1980) models of fuzzy decision analysis, in which the probability laws are fuzzified by invoking the "extension principle" suggested by Zadeh (1975) for propagating membership in functional mappings. However, the theory of confidence-weighted probabilities diverges from fuzzy decision analysis in its treatment of conditionality: The probability of the conditioning event interacts multiplicatively with the Bayes risk/membership function in (4.3′) and (4.5′), which violates the extension principle. This interaction has the important effect of maintaining the piecewise linearity and concavity of all the Bayes risk/membership functions despite the fact that the mapping $\pi \mapsto P_\pi(\mathbf{E}|\mathbf{F})$ is nonlinear if $\mathbf{F}$ is a proper subset of Θ. In fuzzy decision analysis, the laws of lower and upper probabilities are obeyed by level sets of probabilities—that is, sets of probabilities whose degree of membership exceeds a given threshold—*at every level of membership*. In the confidence-weighted probability model, this is generally true only at a membership level of 1.

5. An example of inference with confidence-weighted probabilities. To illustrate the results of the preceding section, consider the assessment given in Table 1 with respect to a three-element set of events, $\Theta = \{\mathbf{E}_1, \mathbf{E}_2, \mathbf{E}_3\}$. For example, 0.3 and 0.6 are asserted to be lower and upper probabilities with confidence 1.0, and 0.4 and 0.5 are asserted to be lower and upper probabilities with confidence 0.5, for the event $\mathbf{E}_1$. This portion of the assessment referring to $\mathbf{E}_1$, denoted $\mathscr{A}_1 = \mathscr{A}_{11} \wedge \mathscr{A}_{12}$, is summarized by the MBR function

TABLE 1
Confidence-weighted probability assessment

Assertion	Event	Lower probability	Upper probability	Confidence weight
$\mathscr{A}_{11}$	$\mathbf{E}_1$	0.3	0.6	1.0
$\mathscr{A}_{12}$	$\mathbf{E}_1$	0.4	0.5	0.5
$\mathscr{A}_{21}$	$\mathbf{E}_2 \mid \overline{\mathbf{E}}_1$	0.2	0.5	1.0
$\mathscr{A}_{22}$	$\mathbf{E}_2 \mid \overline{\mathbf{E}}_1$	0.3	0.4	0.5

$r(\mathbf{E}_1; \mathscr{A}_1)$, whose graph is shown in Figure 3a. Its value at a point x is the minimum of $\mu_x(0.3, 1.0)$, $\mu_x(0.4, 0.5)$, $\mu_{1-x}((1-0.5), 0.5)$ and $\mu_{1-x}((1-0.6), 1.0)$, where μ is the function defined in (4.1) and shown in Figure 2. Extension of this function to the simplex Π via (4.3) yields the Bayes risk function $R(\mathscr{A}_1)$, a contour plot of which is shown in Figure 3b. [The simplex is drawn in triangular coordinates, and contours are given at intervals of 0.05. The shaded area is the set of π on which $R_\pi(\mathscr{A}_1) = 1$, corresponding to the set of x on which $r_x(\mathbf{E}_1; \mathscr{A}_1) = 1$. Note that the contours are parallel lines.] Similarly, the partial assessment referring to $\mathbf{E}_2 \mid \overline{\mathbf{E}}_1$ is summarized by the MBR function $r(\mathbf{E}_2 \mid \overline{\mathbf{E}}_1; \mathscr{A}_2)$ whose graph is shown in Figure 4a; its extension to Π via (4.3′) yields the Bayes risk function $R(\mathscr{A}_2)$ whose contours are shown in Figure 4b. Since this part of the assessment refers to a conditional event, the contours are not parallel. The "cuts" in this case enter the top of the "cake" along lines on which $P_\pi(\mathbf{E}_2 \mid \overline{\mathbf{E}}_1)$ is constant, which radiate from its lower-right vertex.

The entire assessment is summarized by the Bayes risk function $R(\mathscr{A}_{\mathscr{N}})$ whose contours are shown in Figure 5a: This is the pointwise minimum of the two functions in Figures 3b and 4b. (The shaded area in this figure is the set of "medial" distributions enveloped by the greatest lower and least upper probabilities with nonzero confidence under the Koopman—Smith—Good model.) The inferences which can be drawn from this assessment with respect to the unconditional event $\mathbf{E}_2$ are summarized by the MBR function $r(\mathbf{E}_2; \mathscr{A}_{\mathscr{N}})$, whose graph is shown in Figure 5b. This curve is the outline of the shadow which would be cast by the surface in Figure 5a if it were illuminated by *parallel* rays of light traveling horizontally from right to left. Inferences with respect to the conditional event $\mathbf{E}_3 \mid \overline{\mathbf{E}}_2$ are summarized by the MBR function $r(\mathbf{E}_3 \mid \overline{\mathbf{E}}_2; \mathscr{A}_{\mathscr{N}})$, whose graph is shown in Figure 5c. This curve is the outline of the shadow which would be cast by the surface in Figure 5a if it were illuminated by rays of light emanating from a *point source* located at an altitude of 1 unit above the vertex at the top of the figure. The finite sets of informative confidence-weighted probabilities which generate the two inferred MBR functions, as provided in part (ii) of Theorem 3, are listed in Table 2. (These are computed by parametric linear programming. Details of the LP formulation are given in Appendix 2.) Each confidence-weighted lower or upper probability in Table 2 corresponds to one of the dotted lines forming the envelopes of the curves in Figures 5b or 5c.

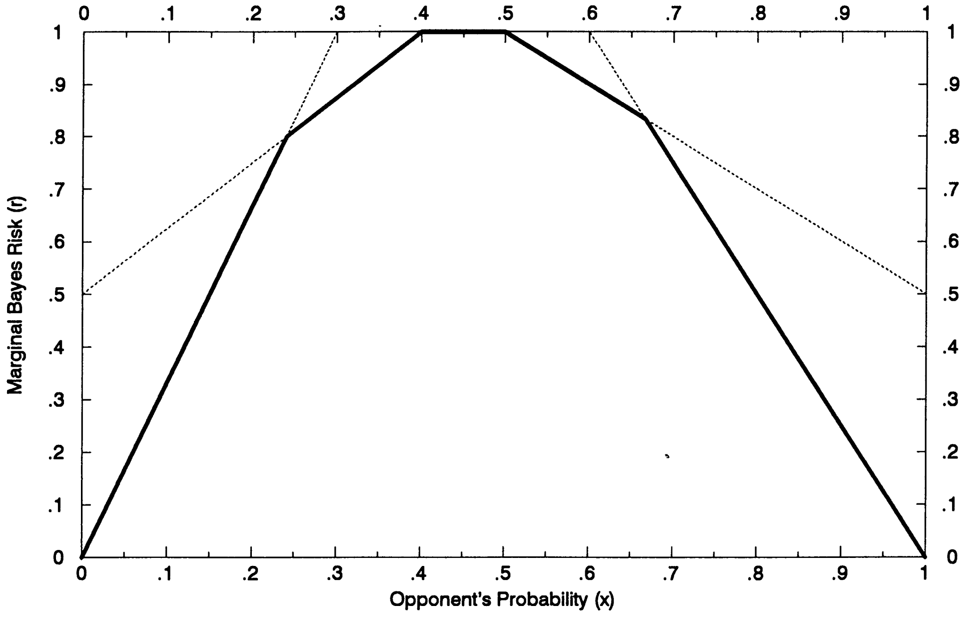

FIG. 3a. *Assessed* MBR *function for the unconditional event* $\mathbf{E}_1$: 0.3 *and* 0.6 *are lower and upper probabilities with confidence* 1.0; 0.4 *and* 0.5 *are lower and upper probabilities with confidence* 0.5.

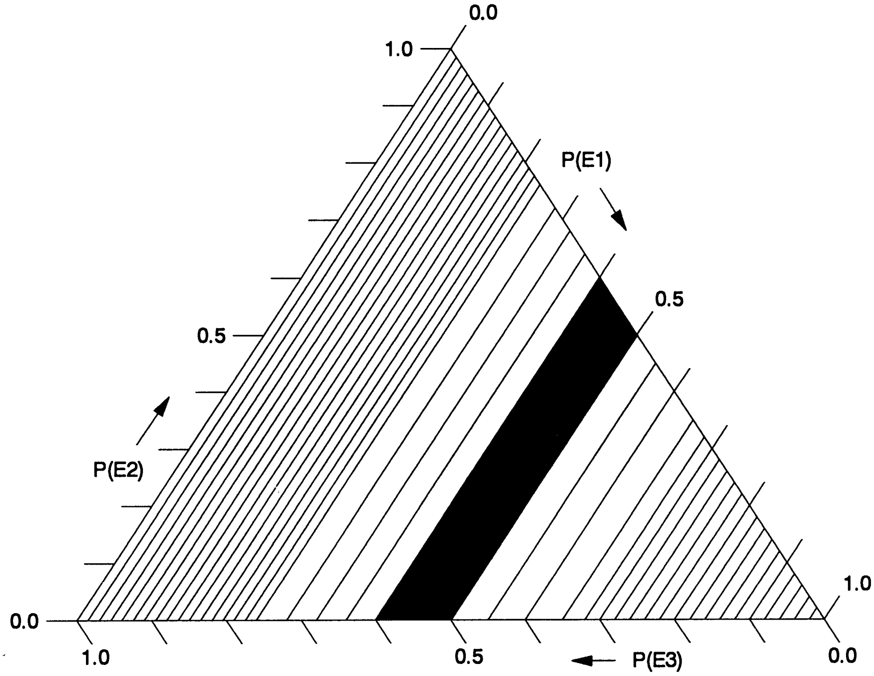

FIG. 3b. *Contour plot of the corresponding Bayes risk function (extension of the* MBR *function in Figure* 3*a to the simplex* Π, *shown in triangular coordinates). Shaded area consists of points at which the Bayes risk equals* 1.0; *surrounding lines are isoquants at heights of* 0.95, 0.90, 0.85, *etc.*

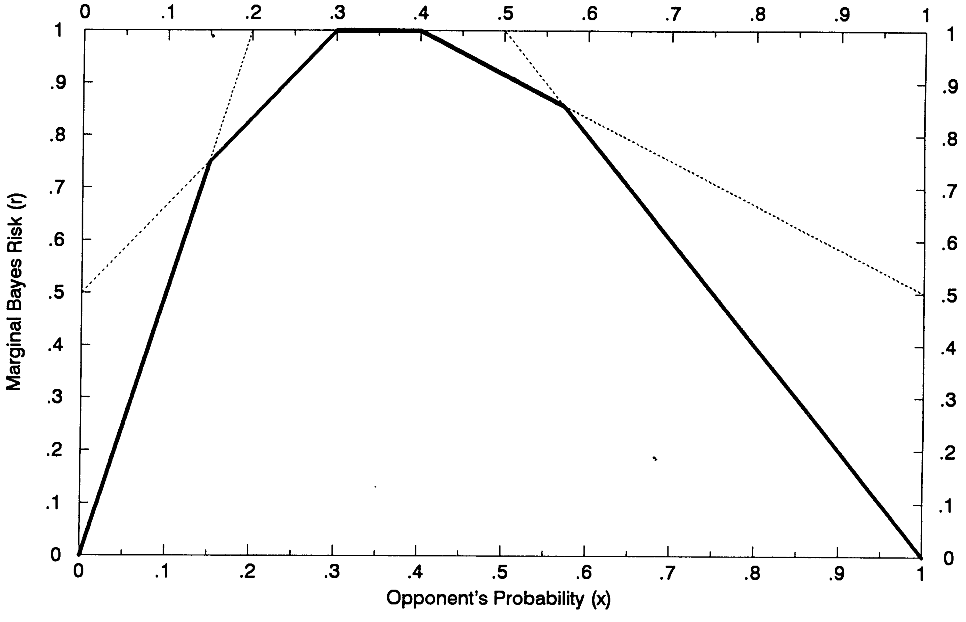

FIG. 4a. *Assessed* MBR *function for the conditional event* $\mathbf{E}_2|\overline{\mathbf{E}}_1$: 0.2 *and* 0.5 *are lower and upper probabilities with confidence* 1.0; 0.3 *and* 0.4 *are lower and upper probabilities with confidence* 0.5.

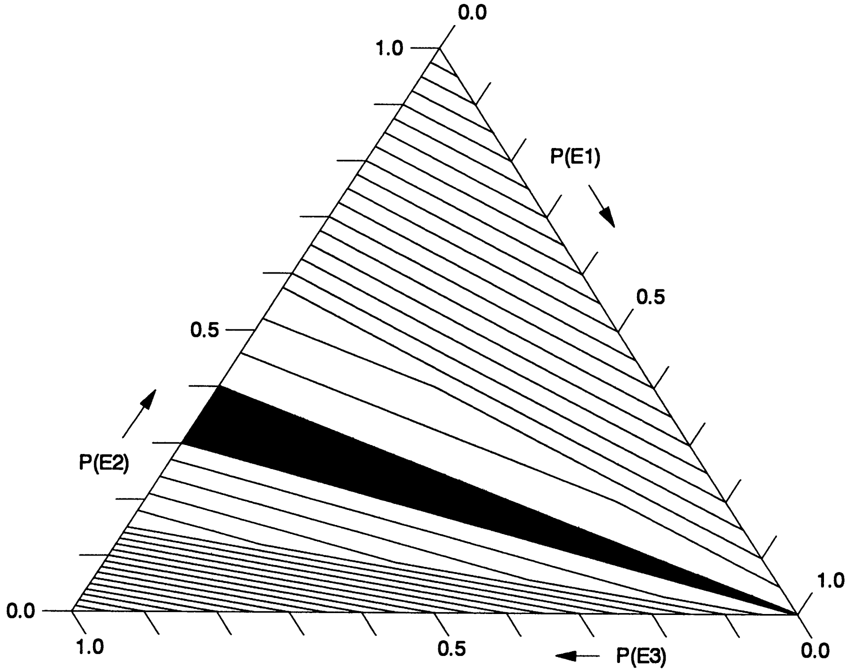

FIG. 4b. *Contour plot of the corresponding Bayes risk function (extension of the* MBR *function in Figure 4a to the simplex* Π). *Isoquants are nonparallel because of conditionality: The probability of the conditioning event* ($\overline{\mathbf{E}}_1$) *goes to* 0 *at the lower-right vertex.*

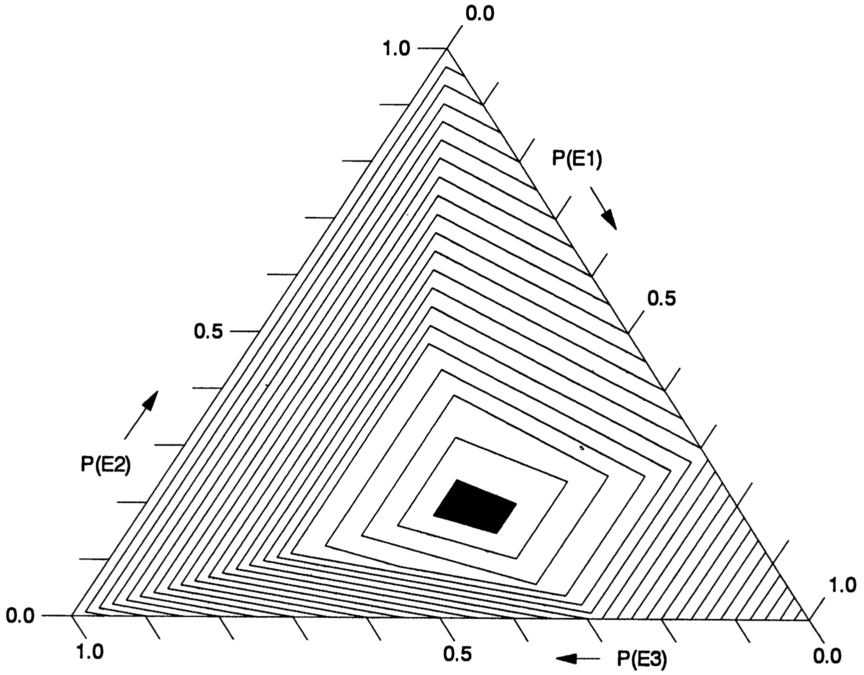

FIG. 5a. *Contour plot of the Bayes risk function induced by the entire assessment (pointwise minimum of the functions plotted in Figures 3b and 4b).*

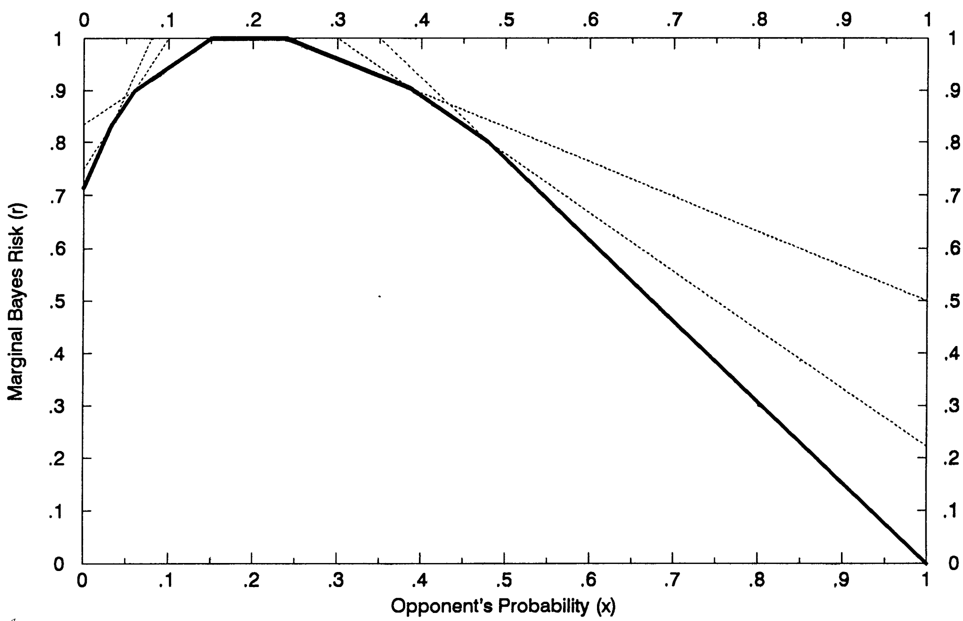

FIG. 5b. *Inferred* MBR *function for the unconditional event* $\mathbf{E}_2$ *(outline of the shadow cast by the surface in Figure 5a when illuminated by parallel rays of light traveling right to left). Confidence-weighed probabilities summarizing this function (corresponding to dotted lines) are given in Table 2.*

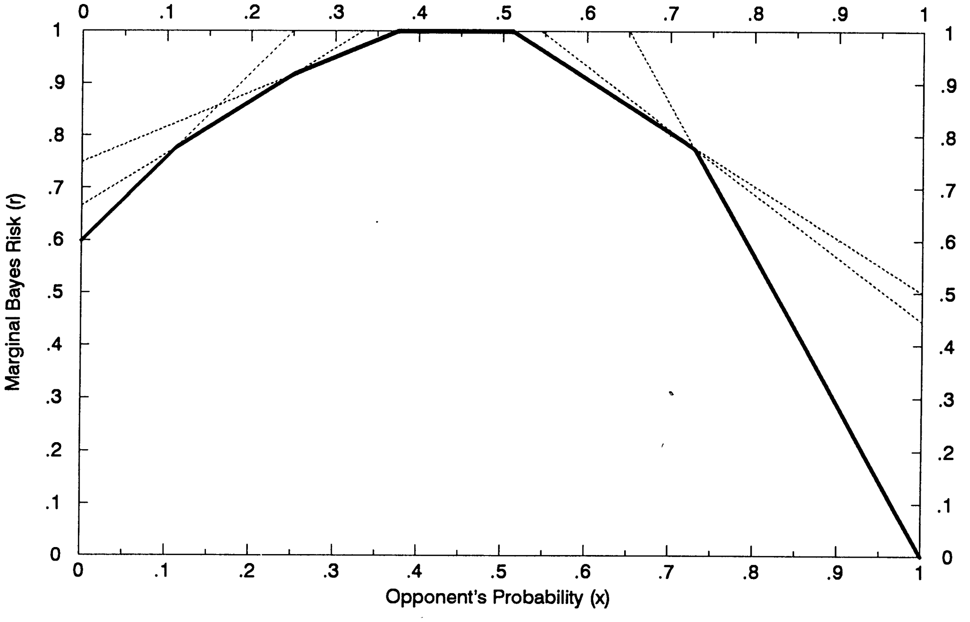

FIG. 5c. *Inferred* MBR *function for the conditional event* $\mathbf{E}_3|\overline{\mathbf{E}}_2$ *(outline of the shadow cast by the surface in Figure* 5a *when illuminated by a point source of light at a height of* 1.0 *unit above the top vertex of the simplex, at which the conditioning event has probability* 0). *Confidence-weighted probabilities summarizing this function (corresponding to dotted lines) are given in Table* 2.

TABLE 2
Inferences from assessment

Event	Lower probability	Upper probability	Confidence weight
$\mathbf{E}_2$	0.08	—	0.2857
$\mathbf{E}_2$	0.10	—	0.2500
$\mathbf{E}_2$	0.15	—	0.1667
$\mathbf{E}_2$	—	0.24	0.5000
$\mathbf{E}_2$	—	0.30	0.7778
$\mathbf{E}_2$	—	0.25	1.0000
$\mathbf{E}_3|\overline{\mathbf{E}}_2$	0.2500	—	0.4000
$\mathbf{E}_3|\overline{\mathbf{E}}_2$	0.3333	—	0.3333
$\mathbf{E}_3|\overline{\mathbf{E}}_2$	0.3750	—	0.2500
$\mathbf{E}_3|\overline{\mathbf{E}}_2$	—	0.5122	0.5000
$\mathbf{E}_3|\overline{\mathbf{E}}_2$	—	0.5455	0.5556
$\mathbf{E}_3|\overline{\mathbf{E}}_2$	—	0.6512	1.0000

Notice that no nontrivial lower probability can be inferred for either target event with a confidence greater than 0.4: The inferences obtained for the target events are less "sharp" than the direct assessments given for the source events. This is a manifestation of the dissipation-of-confidence effect discussed in Section 2. It calls into question the usual divide-and-conquer method of decision analysis, in which complex, hard-to-think-about events are first decomposed into simpler constituents (often by conditioning), probabilities are directly elicited for the latter, and seemingly precise judgments are then constructed for the original events through the application of probability calculus. Under the confidence-weighted probability model, indirect judgments are not necessarily perfect substitutes for direct ones *even in principle*, and the subject is therefore encouraged to assess his uncertainty about events in Θ from many angles, top-down as well as bottom-up.

It is possible in this framework for every conditional or unconditional event defined on Θ to have a lower and upper probability which coincide and which both have confidence 1. (If p is simultaneously a lower and upper probability for $\mathbf{E}$, both with confidence 1, the MBR function for $\mathbf{E}$ is just a triangle with its apex at $x = p$.) In this case, the confidence weights and lower/upper distinctions are superfluous: There is effectively a "determinate" distribution π^0 on Θ consisting of the unique probabilities assigned to its atoms. This situation arises if *and only if* the subject directly assesses a lower probability with confidence 1 for every atom of Θ, and these sum to 1; a decomposed assessment cannot yield this result. (The set of acceptable gambles in this case is the intersection of the orthant of gambles whose minimum element is greater than or equal to -1 with the half-space of gambles whose expectation is nonnegative with respect to π^0. The vertices of this set are precisely the acceptable gambles generated by assigning a lower probability of π^0_m with confidence 1 to the mth atom of Θ for $m \in \{1, \dots, M\}$.) For example, with respect to the three-atom state space discussed above, the subject might assert that 0.45, 0.35 and 0.2 are lower probabilities for $\mathbf{E}_1$, $\mathbf{E}_2$ and $\mathbf{E}_3$, respectively, all with confidence 1. The Bayes risk function on Π for this assessment would be a pyramid with its apex at $\pi^0 = (0.45, 0.35, 0.2)$, whose contours are plotted in Figure 6. The extent to which the subject's actual beliefs are determinate or indeterminate is revealed by the qualitative resemblance of his Bayes risk function to that of Figure 6 or that of Figure 5a.

6. Applications and extensions. These results show that a relaxation of standard axioms of subjective probability, beyond abandonment of completeness, does indeed lead to a nontrivial model of second-order indeterminacy, as many authors have previously conjectured. The model is applicable to the reconciliation of inconsistent probability judgments and to the sensitivity analysis of Bayesian decision models. These applications will now be briefly sketched; more details are given in Nau (1989).

Subjective probability judgments obtained from real decision makers are often inconsistent [Lindley, Tversky and Brown (1979) and Moskowitz and Sarin (1983)], particularly when the source events are conditional. In such

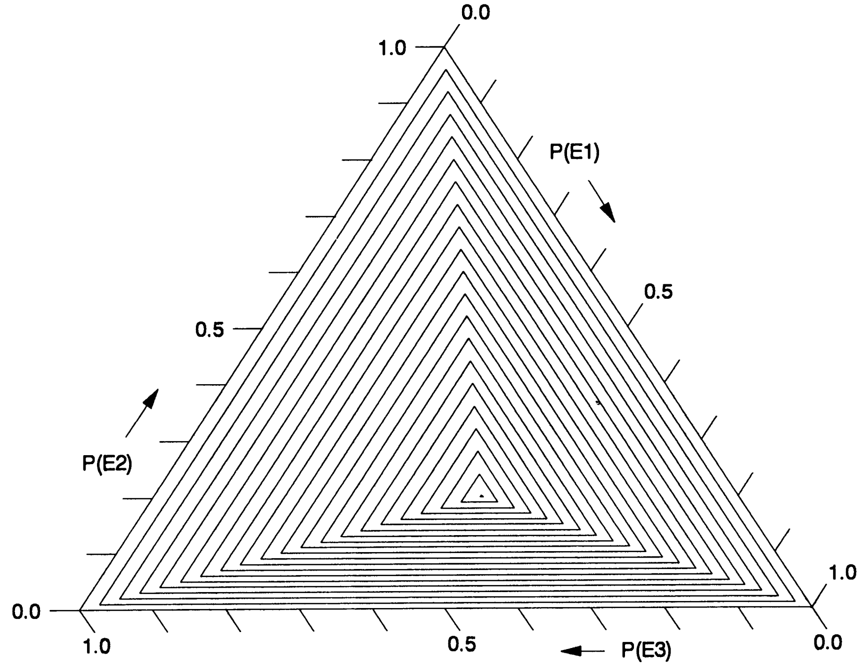

FIG. 6. *Contour plot of the Bayes risk function induced by a "determinate" assessment*: 0.45, 0.35 *and* 0.2 *are lower probabilities with confidence* 1 *for* $\mathbf{E}_1$, $\mathbf{E}_2$ *and* $\mathbf{E}_3$, *respectively.*

cases it is helpful to be able to suggest directions and magnitudes for revisions of the conflicting probability judgments, which entails making tradeoffs: Those which are the most "sure" presumably should be revised the least. The association of a confidence weight with every lower or upper probability provides a basis for such tradeoffs, and an assessment of confidence-weighted probabilities is therefore effectively self-reconciling. The Bayes risk function of an inconsistent (incoherent) assessment achieves a maximum value of less than unity [Theorem 2(i)], but otherwise it qualitatively resembles the Bayes risk function of a consistent assessment: It is not vacuous. The point at which the Bayes risk is maximized provides an anchor for a reconciled assessment. This is the distribution which minimizes the maximum of the terms:

$$c_{nj}\left[P_\pi(\mathbf{F}_n) - P_\pi(\mathbf{E}_n\mathbf{F}_n)/p_{nj}\right] \quad \text{and} \quad c_{nj}\left[P_\pi(\mathbf{F}_n) - P_\pi(\overline{\mathbf{E}}_n\mathbf{F}_n)/(1 - q_{nj})\right]$$

over all $n \in \mathscr{N}$, $j \in \mathscr{J}_n$. [See Appendix 2, (A.2.4).] Note that the assessment of p_{nj} as a lower probability for $\mathbf{E}_{nj}|\mathbf{F}_{nj}$ is equivalent to the constraint $p_{nj} \leq P_\pi(\mathbf{E}_{nj}|\mathbf{F}_{nj})$, whose linearized form is $P_\pi(\mathbf{F}_n) - P_\pi(\mathbf{E}_n\mathbf{F}_n)/p_{nj} \leq 0$, and similarly for upper probabilities. The terms above are deviations from these constraints, scaled by the corresponding confidence weights, and their maximum is nonpositive if and only if the assessment is coherent. Hence, the confidence-weighted probability model supports a minimax-weighted-deviation approach to the reconciliation of incoherence.

As an illustration, suppose the subject whose assessment was given in Table 1 adds that he is quite confident the probability of $\mathbf{E}_3|\overline{\mathbf{E}}_2$ is not more than 0.3 —that is, he adds the judgment that 0.3 is an upper probability for $\mathbf{E}_3|\overline{\mathbf{E}}_2$ with confidence 1. This is inconsistent with the original assessment, which implies 0.375 to be a lower probability for $\mathbf{E}_2|\overline{\mathbf{E}}_2$ with confidence 0.25. The Bayes risk function for the augmented assessment attains a maximum value of 0.966, and this occurs uniquely at $\pi = (0.527, 0.221, 0.252)$, yielding $P_\pi(\mathbf{E}_1) = 0.527$, $P_\pi(\mathbf{E}_2|\overline{\mathbf{E}}_1) = 0.467$ and $P_\pi(\mathbf{E}_3|\overline{\mathbf{E}}_2) = 0.324$ as a suggested reconciliation of the conflicting judgments. Thus, the directly assessed least upper probabilities of $\mathbf{E}_1$, $\mathbf{E}_2|\overline{\mathbf{E}}_1$ and $\mathbf{E}_3|\overline{\mathbf{E}}_2$ are revised upward by 0.027, 0.067 and 0.024, respectively.

In Bayesian sensitivity analysis and robust inference, the unique prior probability distribution of standard Bayesian analysis is typically replaced by a family of nested sets of distributions, parameterized by some measure of "distance" from a set of reference distributions. (The reference set itself may or may not be a singleton.) For the alternative decisions or estimators under consideration, an evaluation function is defined on sets of distributions—for example, minimum or minimax expected loss. The evaluation function is then used to rate the alternatives on sets of distributions within a specified distance from the reference set. The objective is to find an alternative which performs well across a reasonably broad range of distributions—for example, an alternative which is "within δ of being optimal for every distribution within distance ε" from the reference set. [See, e.g., the concepts of "potentially optimal" and "almost potentially optimal" decisions discussed by Rios Insua (1990) or the classes of "ε-contaminated distributions" and the criterion of "ε-procedure robustness" discussed by Berger (1984).]

The need for such extensions of Bayesian decision analysis and inference methods is widely accepted, yet the rationale for introducing families of sets of distributions is extraneous to the standard theory of subjective probability, and consequently the theory provides no guidelines for their construction (e.g., choosing "metrics" with which to measure distances). The results of this paper bring such procedures within the scope of the subjective theory. An assessment of confidence-weighted probabilities is summarized by a Bayes risk function on the probability simplex Π, and the set of distributions whose Bayes risk exceeds a threshold $1 - \varepsilon$ may be considered as the set of distributions "within distance ε" of the subject's reference set of distributions. With regard to sensitivity analysis, this subjective concept of distance has been applied to finite-state decision models by Nau (1989), yielding results very similar to the minimum-weighted-distance criterion of Fishburn, Murphy and Isaacs (1968) and the almost potentially optimal criterion of Rios Insua (1990).

With regard to Bayesian robustness, the set of distributions whose Bayes risk exceeds $1 - \varepsilon$ in the maximally confident assessment of Figure 6 is precisely the set of ε-contaminated distributions whose contamination class is (as usual) the entire set Π. In other words, the contours of the Bayes risk function in Figure 6 are the boundaries of sets of ε-contaminated distributions centered on the reference distribution π^0. In the more general case exemplified by Figure 5a, the diameter of the set of distributions whose Bayes risk exceeds $1 - \varepsilon$ is a concave rather than a linear function of ε, and the reference set (corresponding to $\varepsilon = 0$) need not be a singleton. Of course, robust Bayesian analysis typically deals with situations for which the elicitation of classes of priors in terms of confidence-weighted probabilities would be impractical: State spaces are infinite, so classes of priors must be characterized economically, and their subjective origins are often downplayed.

The axioms which have been used here explicitly require either risk neutrality or constant absolute risk aversion, and no attempt has been made to separate nonlinear utility from belief indeterminacy. [Indeed, this would be difficult to achieve in practice for reasons noted by Leamer (1986) and Kadane and Winkler (1988.] However, the confidence-weighted-probability model can be formally extended in a straightforward way to deal jointly with probability and utility through the device of extraneous scaling probabilities or "horse lotteries" in the manner of Anscombe and Aumann (1963). This leads to a theory of confidence-weighted subjective expected utilities over finite sets of events and consequences, a generalization of the partially ordered preference theory of Seidenfeld, Kadane and Schervish (1989); details appear in Nau (1990).

APPENDIX 1. Proofs of theorems

For part (i) of Theorem 1, note that the cancellation axiom implies that the truth value of "$\mathbf{X} \succsim \mathbf{Y}$" depends only on $\mathbf{X} - \mathbf{Y}$, whence any preference relation satisfying this axiom is defined by $\mathbf{X} \succsim \mathbf{Y} \Leftrightarrow \mathbf{X} - \mathbf{Y} \in \mathscr{B}$ for *some* subset $\mathscr{B}$ of $\mathbb{R}^M$. If it satisfies A.1 and also $\mathbf{X}_n \succsim \mathbf{Y}_n$ for all $n \in \mathscr{N}$, then $\mathscr{B}$ minimally must contain $\mathbf{0}$ and $\mathbf{X}_n - \mathbf{Y}_n$ for all $n \in \mathscr{N}$. The cancellation and convexity axioms together imply that $\{\mathbf{X} \succsim \mathbf{Y}$ and $\mathbf{X}' \succsim \mathbf{Y}'\} \Rightarrow \alpha\mathbf{X} + (1-\alpha)\mathbf{Y} \succsim \alpha\mathbf{X}' + (1-\alpha)\mathbf{Y}'$, whence $\mathscr{B}$ must be convex. The dominance axiom implies that any vector which dominates a vector in $\mathscr{B}$ is also in $\mathscr{B}$. Taken together, these operations establish that the set $\mathscr{B}$ defining a convex extension of $\mathscr{A}_{\mathscr{N}}$ must contain the set $\mathscr{B}_{\mathscr{N}} = \mathrm{CONV}^+(\mathbf{0}, \mathbf{X}_1 - \mathbf{Y}_1, \dots, \mathbf{X}_N - \mathbf{Y}_N\}$, but no more than this; hence, the preference relation defined by $\mathscr{B}_{\mathscr{N}}$ is the minimal convex extension. For part (ii), note that every assessment has a minimal convex extension as defined above; this is coherent if and only if there do not exist $\mathbf{X}$ and $\mathbf{Y}$ such that $\mathbf{X} \succsim \mathbf{Y}$ and $\mathbf{X} >^* \mathbf{Y}$, which by part (i) is equivalent to the condition $\mathbf{X} - \mathbf{Y} \in \mathscr{B}_{\mathscr{N}}$ and $\mathbf{X} - \mathbf{Y} < 0$.

The remaining proofs invoke a familiar separating-hyperplane theorem of linear algebra [Gale (1960), Theorem 2.8] stating that exactly one of the following systems has a solution:

$$\text{(I)} \qquad \mathbf{A}\mathbf{x} \le \mathbf{b}, \qquad \mathbf{x} \ge 0,$$

$$\text{(II)} \qquad \mathbf{w}^{\mathrm{T}}\mathbf{A} \ge 0, \qquad \mathbf{w}^{\mathrm{T}}\mathbf{b} < 0, \qquad \mathbf{w} \ge 0,$$

where $\mathbf{A}$ is an arbitrary matrix, and "T" denotes transposition. Let $\mathbf{A}$ be the matrix whose nth column is the vector $(1, c_n\mathbf{B}_n^*)$ for $n = 1, \dots, N$; and let

$\mathbf{b} = (1, \mathbf{B})$ and $\mathbf{w} = (z, \pi)$, where $\mathbf{B}$ and π are M vectors and z is a scalar. Note that any solution to (II) must satisfy $\Sigma_{\theta \in \Theta}\pi(\theta) > 0$, and hence may be normalized so that $\Sigma_{\theta \in \Theta}\pi(\theta) = 1$. Therefore, if (II) has a solution it may be assumed to satisfy $\pi \in \Pi$. With these values for $\mathbf{A}$, $\mathbf{b}$ and $\mathbf{w}$, (I) has a solution if and only if $\mathbf{B} \in \mathscr{B}_{\mathscr{N}}$, and (II) has a solution if and only if there exists a $\pi \in \Pi$ and $z \geq 0$ such that $1 - c_n P_\pi(-\mathbf{B}_n^*) \geq 1 - z > 1 - P_\pi(-\mathbf{B})$ for all $n \in N$, which is equivalent to $R_\pi(\mathscr{A}_{\mathscr{N}}) > 1 + \min\{0, P_\pi(\mathbf{B})\}$. Hence, by the separating-hyperplane theorem, $\mathbf{B} \in \mathscr{B}_{\mathscr{N}}$ if and only if $R_\pi(\mathscr{A}_{\mathscr{N}}) \leq 1 + \min\{0, P_\pi(\mathbf{B})\}$ for all $\pi \in \Pi$. Part (i) of Theorem 2 then follows by letting $\mathbf{B} = -\varepsilon\mathbf{1}$ as $\varepsilon \to 0^+$. Part (ii) follows by letting $\mathbf{B} = c(\mathbf{X} - \mathbf{Y})^*$.

For part (iii) of Theorem 3, let $\mathbf{B} = (c/p)(\mathbf{E} - p\mathbf{1})\mathbf{F}$. Then $\mathbf{B} \in \mathscr{B}_{\mathscr{N}}$ (i.e., $\mathscr{A}$ is inferable from $\mathscr{A}_{\mathscr{N}}$) if and only if

$$\begin{aligned} 1 - R_\pi(\mathscr{A}_{\mathscr{N}}) &\geq -(c/p)P_\pi((\mathbf{E} - p\mathbf{1})\mathbf{F}) \\ &= (c/p)[pP_\pi(\mathbf{F}) - P_\pi(\mathbf{EF})] \qquad \forall\, \pi \in \Pi. \end{aligned}$$

Since $R_\pi(\mathscr{A}_{\mathscr{N}}) \leq 1$ on Π, this holds trivially wherever $P_\pi(\mathbf{F}) = P_\pi(\mathbf{EF}) = 0$. Elsewhere, it is equivalent to

$$\begin{aligned} &[1 - R_\pi(\mathscr{A}_{\mathscr{N}})]/P_\pi(\mathbf{F}) \geq (c/p)[p - P_\pi(\mathbf{E}|\mathbf{F})] \ \forall\, \pi \in \Pi\colon P_\pi(\mathbf{F}) > 0 \\ &\Leftrightarrow 1 - [1 - R_\pi(\mathscr{A}_{\mathscr{N}})]/P_\pi(\mathbf{F}) \leq 1 - (c/p)[p - P_\pi(\mathbf{E}|\mathbf{F})] \\ &\qquad\qquad \forall\, \pi \in \Pi\colon P_\pi(\mathbf{F}) > 0 \\ &\Leftrightarrow \sup_{\pi\colon P_\pi(\mathbf{E}|\mathbf{F}) = x,\, P_\pi(\mathbf{F}) > 0} \{1 - [1 - R_\pi(\mathscr{A}_{\mathscr{N}})]/P_\pi(\mathbf{F})\} \leq 1 - (c/p)[p - x] \\ &\qquad\qquad \forall\, x \in [0,1] \\ &\Leftrightarrow r_x(\mathbf{E}|\mathbf{F}; \mathscr{A}_{\mathscr{N}}) \leq 1 + (c/p)[x - p] \ \forall\, x \in [0,1]. \end{aligned}$$

Since $R_\pi(\mathscr{A}_{\mathscr{N}}) \leq 1$, it follows that $r_x(\mathbf{E}|\mathbf{F}; \mathscr{A}_{\mathscr{N}}) \leq 1$, whence the last inequality can be rewritten w.l.o.g. as

$$r_x(\mathbf{E}|\mathbf{F}; \mathscr{A}_{\mathscr{N}}) \leq 1 + (c/p)\min\{0, x - p\} = \mu_x(p, c) \qquad \forall\, x \in [0,1],$$

as asserted.

For parts (i) and (ii) of Theorem 3, note that $0 \leq R_\pi(\mathscr{A}_{\mathscr{N}}) \leq 1$ for all $\pi \in \Pi$ implies $0 \leq r_x(\mathbf{E}|\mathbf{F}; \mathscr{A}_{\mathscr{N}}) \leq 1$ for all $x \in [0,1]$; and if $R_\pi(\mathscr{A}_{\mathscr{N}}) = 1$ for some π satisfying $P_\pi(\mathbf{F}) > 0$, then $r_x(\mathbf{E}|\mathbf{F}; \mathscr{A}_{\mathscr{N}}) = 1$ at $x = P_\pi(\mathbf{E}|\mathbf{F})$. To establish concavity and piecewise linearity, it remains to show that $r(\mathbf{E}|\mathbf{F}; \mathscr{A}_{\mathscr{N}})$ is a polyhedral function, or equivalently, that the surgraph of $r(\mathbf{E}|\mathbf{F}; \mathscr{A}_{\mathscr{N}})$—that is, the set of all (x, y) such that $y \leq r_x(\mathbf{E}|\mathbf{F}; \mathscr{A}_{\mathscr{N}})$—is a polyhedral convex set. To show this, note that the surgraph of $r(\mathbf{E}|\mathbf{F}; \mathscr{A}_{\mathscr{N}})$ on $[0,1]$ is the image of the surgraph of $R(\mathscr{A}_{\mathscr{N}})$ on Π under the mapping $(\pi, z) \mapsto (x, y)$ defined by

$$\text{(A.1.1)} \qquad x = P_\pi(\mathbf{E}|\mathbf{F}), \qquad y = 1 - (1 - z)/P_\pi(\mathbf{F}).$$

While this transformation is not linear, it is *linearity preserving* in the sense that if (π, z) and (π', z') are mapped into (x, y) and (x', y'), respectively, then every convex combination of (π, z) and (π', z') is mapped into a convex combination of (x, y) and (x', y'). In particular, $\alpha(\pi, z) + (1 - \alpha)(\pi', z')$ is mapped into $\beta(x, y) + (1 - \beta)(x', y')$, where

$$\beta = \frac{\alpha P_\pi(\mathbf{F})}{\alpha P_\pi(\mathbf{F}) + (1 - \alpha) P_{\pi'}(\mathbf{F})},$$

and this relation between α and β is a one-to-one mapping of [0, 1] onto [0, 1] provided $P_\pi(\mathbf{F}) > 0$ and $P_{\pi'}(\mathbf{F}) > 0$. However, points in Π at which $\mathbf{F}$ has zero probability can be ignored during the construction of $r(\mathbf{E}|\mathbf{F}; \mathscr{A}_{\mathscr{N}})$, because the quantity

$$\max_{\pi:\, P_\pi(\mathbf{E}|\mathbf{F}) = x,\, P_\pi(\mathbf{F}) = \varepsilon} 1 - \frac{1 - R_\pi(\mathscr{A}_{\mathscr{N}})}{P_\pi(\mathbf{F})}$$

either remains constant or goes to $-\infty$ as $\varepsilon \to 0+$ while x remains fixed, by virtue of the local linearity of $R(\mathscr{A}_{\mathscr{N}})$. That is, for some positive ε, we have for all $x \in [0, 1]$:

$$r_x(\mathbf{E}|\mathbf{F}; \mathscr{A}_{\mathscr{N}}) = \max_{\pi:\, P_\pi(\mathbf{E}|\mathbf{F}) = x,\, P_\pi(\mathbf{F}) \geq \varepsilon} 1 - \frac{1 - R_\pi(\mathscr{A}_{\mathscr{N}})}{P_\pi(\mathbf{F})}.$$

It follows that the surgraph of $r(\mathbf{E}|\mathbf{F}; \mathscr{A}_{\mathscr{N}})$ is the convex hull of the images, under the mapping defined by (A.1.1), of the extreme points of the surgraph of $R(\mathscr{A}_{\mathscr{N}})$ at which $\mathbf{F}$ has positive probability. Hence, $r(\mathbf{E}|\mathbf{F}; \mathscr{A}_{\mathscr{N}})$ inherits the polyhedrality of $R(\mathscr{A}_{\mathscr{N}})$, as asserted. □

APPENDIX 2. Computation of confidence-weighted probabilities

Theorem 3 states that the inferences which can be derived from an assessment $\mathscr{A}_{\mathscr{N}}$ with respect to a target event $\mathbf{E}|\mathbf{F}$ are summarized by a piecewise linear MBR function, $r(\mathbf{E}|\mathbf{F}; \mathscr{A}_{\mathscr{N}})$, which is generated by finite sets of confidence-weighted lower and upper probabilities. (For example, see Figures 5b and 5c and Table 2.) These can be computed in practice by linear programming in the following way. The graph of the MBR function is a polygon, so it is defined by its vertices, which can be found as points of tangency with lines of differing slopes. For example, for any $\tau \in [-1, 1]$, the point of tangency with the line having slope $-\tau/(1 - |\tau|)$ can be found by solving

$$\min_{x \in [0, 1]} \tau x - (1 - |\tau|) r_x(\mathbf{E}|\mathbf{F}; \mathscr{A}_{\mathscr{N}}).$$

Using (4.5′), this can be rewritten as a minimization over π:

$$\min_{\pi \in \Pi:\, P_\pi(\mathbf{F}) > \varepsilon} \tau P_\pi(\mathbf{E}|\mathbf{F}) - (1 - |\tau|)\left[1 - (1 - R_\pi(\mathscr{A}_{\mathscr{N}}))/P_\pi(\mathbf{F})\right]$$

for some small $\varepsilon > 0$. Dropping the constant term and expanding $P_\pi(\mathbf{E}|\mathbf{F})$ yields

$$\text{(A.2.1)} \qquad \min_{\pi \in \Pi:\, P_\pi(\mathbf{F}) \geq \varepsilon} \left[\tau P_\pi(\mathbf{EF}) + (1 - |\tau|)(1 - R_\pi(\mathscr{A}_{\mathscr{N}}))\right]/P_\pi(\mathbf{F}).$$

The presence of $P_\pi(\mathbf{F})$ in the denominator renders this problem nonlinear, but it can be linearized by solving

$$\text{(A.2.2)} \qquad \min_{\pi \in \Pi:\, P_\pi(\mathbf{F}) \geq \varepsilon} \tau P_\pi(\mathbf{EF}) + (1 - |\tau|)(1 - R_\pi(\mathscr{A}_{\mathscr{N}})) - \lambda P_\pi(\mathbf{F})$$

parametrically in λ, and finding the largest λ that yields an optimal objective value not less than 0. The value of π that solves (A.2.2) for this λ is also a solution to (A.2.1). In turn, (A.2.2) is equivalent to the following linear program:

$$\text{(A.2.3)} \qquad \begin{gathered} \min_{(\pi, z) \in \Pi \times [0,1],\, P_\pi(\mathbf{F}) \geq \varepsilon} \tau P_\pi(\mathbf{EF}) + (1 - |\tau|) z - \lambda P_\pi(\mathbf{F}) \\ \text{subject to } z \geq 1 - R_\pi(\mathscr{A}_{\mathscr{N}}), \end{gathered}$$

in which a new scalar variable z has been introduced. The constraint $z \geq 1 - R_\pi(\mathscr{A}_{\mathscr{N}})$ represents the system of inequalities:

$$\text{(A.2.4)} \qquad \begin{aligned} &z \geq c_{nj}\left[P_\pi(\mathbf{F}_n) - P_\pi(\mathbf{E}_n \mathbf{F}_n)/p_{nj}\right], \\ &z \geq c_{nj}\left[P_\pi(\mathbf{F}_n) - P_\pi(\overline{\mathbf{E}}_n \mathbf{F}_n)/(1 - q_{nj})\right] \qquad \forall\, n \in \mathscr{N},\ j \in \mathscr{J}_n, \end{aligned}$$

together with the implicit constraint $z \geq 0$. [Compare with (4.3′) and (4.4), noting that $P_\pi(\mathbf{F}_n) x_n(\pi) = P_\pi(\mathbf{F}_n) P_\pi(\mathbf{E}_n | \mathbf{F}_n) = P_\pi(\mathbf{E}_n \mathbf{F}_n)$, etc.] The recipe for constructing $r(\mathbf{E}|\mathbf{F}; \mathscr{A}_{\mathscr{N}})$ is as follows: For some value of τ in the interval $[-1, 1]$, solve (A.2.3) parametrically in λ and find the largest λ for which the optimal objective value is not less than 0. Let $\hat{\pi}$ and $\hat{z}$ denote the optimal values of π and z for this λ, let $\hat{x} = P_{\hat{\pi}}(\mathbf{E}|\mathbf{F})$, and let $\hat{y} = 1 - \hat{z}/P_{\hat{\pi}}(\mathbf{F})$. Then $(\hat{x}, \hat{y})$ is a vertex on the graph of $r_x(\mathbf{E}|\mathbf{F}; \mathscr{A}_{\mathscr{N}})$ versus x that is touched by the tangent line with slope $-\tau/(1 - |\tau|)$. By a finite search procedure, this process can be repeated for different values of τ until all the vertices have been found. When lines are extrapolated through pairs of adjacent vertices, each line defines one of the confidence-weighted lower or upper probabilities $(\hat{p}_s, \hat{c}_s)$ or $(\hat{q}_t, \hat{c}_t)$ generating $r(\mathbf{E}|\mathbf{F}; \mathscr{A}_{\mathscr{N}})$. (For example, $\hat{p}_s$ is the $y = 1$ intercept and $\hat{c}_s$ is 1 minus the $x = 0$ intercept of the sth such line with positive slope.) If $\mathbf{F}$ is logically certain, the λ term can be dropped from the objective function, and the linear program can simply be solved parametrically in τ over the interval $[-1, 1]$. In this case, vertices will be found only at values of τ where a change of basis occurs. Software which implements this algorithm and plots graphs such as those in Figures 3, 4 and 5 is available from the author on request. The software is written in STSC APL* Plus/PC.

Acknowledgments. The author is grateful to Peter Wakker, Robert Winkler, an Associate Editor and several anonymous referees for their comments on earlier drafts, and to Dennis Lindley, Howard Kunreuther and Larry Wasserman for helpful discussions. The opinions expressed herein and any remaining errors are the sole responsibility of the the author.

REFERENCES

ANSCOMBE, F. and AUMANN, R. (1963). A definition of subjective probability. *Ann. Math. Statist.* **34** 199–205.

AUMANN, R. (1962). Utility theory without the completeness axiom. *Econometrica* **30** 445–462.

BELL, D. E. (1982). Regret in decision making under uncertainty. *Oper. Res.* **30** 961–981.

BERGER, J. O. (1984). The robust Bayesian viewpoint (with discussion). In *Robustness of Bayesian Analyses* (J. B. Kadane, ed.) 63–144. North-Holland, Amsterdam.

BERGER, J. and BERLINER, L. M. (1986). Robust Bayes and empirical Bayes analysis with ε-contaminated priors. *Ann. Statist.* **14** 461–486.

BUEHLER, R. J. (1976). Coherent preferences. *Ann. Statist.* **4** 1051–1064.

DE FINETTI, B. (1937). La prévision: Ses lois logiques, ses sources subjectives. *Ann. Inst. H. Poincaré Sect. B (N.S.)* **7** 1–68. Translation reprinted in Kyburg and Smokler (1980).

DE FINETTI, B. (1974). *Theory of Probability* **1**. Wiley, New York.

DEGROOT, M. (1970). *Optimal Statistical Decisions*. McGraw-Hill, New York.

DEROBERTIS, L. and HARTIGAN, J. A. (1981). Bayesian inference using intervals of measures. *Ann. Statist.* **9** 235–244.

DUBOIS, D. and PRADE, H. (1989). Fuzzy sets, probability, and measurement. *European J. Oper. Res.* **40** 135–154.

FISHBURN, P. C. (1983a). Ellsberg revisited: A new look at comparative probability. *Ann. Statist.* **11** 1047–1059.

FISHBURN, P. C. (1983b). A generalization of comparative probability on finite sets. *J. Math. Psych.* **27** 298–310.

FISHBURN, P. C., MURPHY, A. H. and ISAACS, H. H. (1968). Sensitivity of decisions to probability estimation errors: A reexamination. *Oper. Res.* **16** 254–267.

FREELING, A. N. S. (1980). Fuzzy sets and decision analysis. *IEEE Trans. Systems Man Cybernet.* **SMC-10** 341–354.

GALE, D. (1960). *The Theory of Linear Economic Models*. McGraw-Hill, New York.

GÄRDENFORS, P. and SAHLIN, N. (1982). Unreliable probabilities, risk-taking, and decision making. *Synthese* **53** 361–386.

GÄRDENFORS, P. and SAHLIN, N. (1983). Decision making with unreliable probabilities. *British J. Math. Statist. Psych.* **36** 240–251.

GOOD, I. J. (1962). Subjective probability as the measure of a non-measurable set. In *Logic, Methodology, and Philosophy of Science* (E. Nagel, P. Suppes and A. Tarski, eds.) 319–339. Stanford Univ. Press. Reprinted in Kyburg and Smokler (1980).

KADANE, J. B. and WINKLER, R. L. (1988). Separating probability elicitation from utilities. *J. Amer. Statist. Assoc.* **83** 357–363.

KOOPMAN, B. O. (1940). Axioms and algebra of intuitive probability. *Ann. Math.* **41** 269–292.

KYBURG, H. E. (1974). *The Logical Foundations of Statistical Inference*. Reidel, Dordrecht.

KYBURG, H. E. and SMOKLER, H. E. (1980). *Studies in Subjective Probability*, 2nd ed. Krieger, New York.

LEAMER, E. A. (1986). Bid-ask spreads for subjective probabilities. In *Bayesian Inference and Decision Techniques* (P. Goel and A. Zellner, eds.) 217–232. Elsevier, New York.

LEVI, I. (1980). *The Enterprise of Knowledge*. MIT Press.

LEVI, I. (1985). Imprecision and indeterminacy in probability judgment. *Philos. Sci.* **52** 390–409.

LINDLEY, D. V., TVERSKY, A. and BROWN, R. V. (1979). On the reconciliation of probability assessments (with discussion). *J. Roy. Statist. Soc. Ser. A* **142** 146–180.

LOOMES, G. and SUGDEN, R. (1982). Regret theory: An alternative theory of rational choice under uncertainty. *Economic Journal* **92** 805–824.

LUCE, R. D. and NARENS, L. (1985). Classification of concatenation measurement structures according to scale type. *J. Math. Psych.* **29** 1–72.

MARSCHAK, J., BORCH, K., CHERNOFF, H., DEGROOT, M., DORFMAN, R., EDWARDS, W., FERGUSON, T., MIYASAWA, K., RANDOLPH, P., SAVAGE, L., SCHLAIFER, R. and WINKLER, R. (1975). Personal probabilities of probabilities. *Theory and Decision* **6** 121–153.

MOSKOWITZ, H. and SARIN, R. (1983). Improving the consistency of conditional probability assessments for forecasting and decision making. *Management Sci.* **29** 735–749.

NAU, R. (1989). Decision analysis with indeterminate or incoherent probabilities. *Ann. Oper. Res.* **19** 375–403.

NAU, R. (1990). Indeterminate probabilities and utilities on finite sets. Working paper 8929, Fuqua School of Business, Duke Univ.

RIOS INSUA, D. (1990). *Sensitivity Analysis in Multi-Objective Decision Analysis*. Springer, Berlin.

ROCKAFELLAR, R. T. (1970). *Convex Analysis*. Princeton Univ. Press.

SAVAGE, L. J. (1972). *The Foundations of Statistics*, 2nd rev. ed. Dover, New York.

SEIDENFELD, T., KADANE, J. B. and SCHERVISH, M. J. (1989). A representation of partially ordered preferences, Technical report 453, Dept. Statistics, Carnegie Mellon Univ.

SEIDENFELD, T., SCHERVISH, M. J. and KADANE, J. B. (1990). Decisions without ordering. In *Acting and Reflecting* (W. Sieg, ed.) 143–170. Kluwer, Dordrecht.

SMITH, C. A. B. (1961). Consistency in statistical inference and decision. *J. Roy. Statist. Soc. Ser. B* **23** 1–25.

SUPPES, P. (1974). The measurement of belief. *J. Roy. Statist. Soc. Ser. B* **36** 160–175.

WAKKER, P. P. (1990). A behavioral foundation for fuzzy measures. *Fuzzy Sets and Systems* **37** 327–350.

WALLEY, P. (1982). The elicitation and aggregation of beliefs. Technical report, Univ. Warwick.

WALLEY, P. (1991). *Statistical Reasoning with Imprecise Probabilities*. Chapman and Hall, London.

WASSERMAN, L. A. (1990). Prior envelopes based on belief functions. *Ann. Statist.* **18** 454–464.

WATSON, S. R., WEISS, J. J. and DONNELL, M. L. (1979). Fuzzy decision analysis. *IEEE Trans. Systems Man Cybernet.* **SMC-9** 1–9.

WILLIAMS, P. M. (1976). Indeterminate probabilities. In *Formal Methods in the Methodology of Empirical Sciences* (M. Przełęcki, K. Szaniawski and R. Wójcicki, eds.) 229–246. Ossolineum and Reidel, Dordrecht.

ZADEH, L. A. (1965). Fuzzy sets. *Inform. and Control* **8** 338–353.

ZADEH, L. A. (1975). The concept of a linguistic variable and its application to approximate reasoning. *Inform. Sci.* **8** 199–249.

THE FUQUA SCHOOL OF BUSINESS
DUKE UNIVERSITY
DURHAM, NORTH CAROLINA 27706