

AD-A250 540



NTATION PAGE

Form Approved
OPM No. 0704-0188

2

page 1 hour per response, including the time for reviewing instructions, searching existing data sources gathering and
and comments regarding this burden estimate or any other aspect of this collection of information, including suggestions
for information Operations and Reports, 1218 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to
Budget, Washington, DC 20503.

1. AGENCY USE ONLY (Leave Blank)		2. REPORT DATE 1989	3. REPORT TYPE AND DATES COVERED Unknown
4. TITLE AND SUBTITLE Interval-Based Decisions for Reasoning Systems			5. FUNDING NUMBERS DAAB10-86-C-0567
6. AUTHOR(S) Ronald Loui Jerome Feldman Henry Kyburg, Jr.			8. PERFORMING ORGANIZATION REPORT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Departments of Comptuer Science and Philosophy University of Rochester Rochester, NY 14627			10. SPONSORING/MONITORING AGENCY REPORT NUMBER 92-TRF-0022
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) U.S. Army CECOM Signals Warfare Vint Hill Farms Station Warrenton, VA 22186-5100			
1. SUPPLEMENTARY NOTES			
12a. DISTRIBUTION/AVAILABILITY STATEMENT Statement A; Approved for public release, distribution unlimited.			12b. DISTRIBUTION CODE
13. ABSTRACT (Maximum 200 words) This essay trie to expound a conception of interval measunes that permits a particular approach to partial ignorance decision problems. The virtue of this approach for artificial reasoning systems is that the following questions become moot: 1. which secondary criterion to the approach is the need forexplicit epistemological foundations: for instance, a rule of acceptance with a parameter that allow various attitudes toward error. Note that epistemological foundations are already disirable for independent reasons.			
14. SUBJECT TERMS artificial intelligence, data fusion, interval-based reasoning mathematical theory of evidence, bayesian interence			15. NUMBER OF PAGES 8
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED			16. PRICE CODE
18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED		19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL

DTIC
ELECTE
MAY 26 1992
S A D

Interval-Based Decisions for Reasoning Systems

Ronald P. Loui

Jerome A. Feldman

Henry G. E. Kyburg, Jr.

Departments of Computer Science and Philosophy
University of Rochester
Rochester, NY 14627

Abstract and Claims.

This essay tries to expound a conception of interval measures that permits a particular approach to partial ignorance decision problems. The virtue of this approach for artificial reasoning systems is that the following questions become moot: 1. which secondary criterion to apply after maximizing expected utility, and 2. how much indeterminacy to represent. The cost of the approach is the need for explicit epistemological foundations: for instance, a *rule of acceptance* with a parameter that allows various attitudes toward error. Note that epistemological foundations are already desirable for independent reasons.

The development is as follows: 1. probability intervals are useful and natural in A.I. systems; 2. wide intervals avoid error, but are useless in some risk-sensitive decision-making; 3. yet one may obtain narrower, or otherwise decisive intervals with a more relaxed attitude toward error; 4. if bodies of knowledge can be ordered by their attitude to error, one should perform the decision analysis with the acceptable body of knowledge that allows the least error, of those that are useful. The resulting behavior differs from that of a Bayesian probabilist because in the proposal, 5. intervals based on successive bodies of knowledge are not always nested; 6. the use of a probability for a particular decision does not require commitment to the probability for credence; and 7. there may be no acceptable body of knowledge that is useful; hence, sometimes no decision is mandated.

1. Interval Measures.

By now, the use of an interval measure is regarded highly for probability judgements in reasoning systems. Researchers selecting formalisms for quantifying belief have all recognized the virtues of (partial)¹ indeterminacy in probability judgement ([Bar81], [GLF81], [Dil82], [Low82], [WeH82], [Qui83], [Wes83], [Gin84], [Lus84], [Str84], etc.).

Intervals allow varying degrees of commitment in probability assertion. At the extremes, $P(A) = [0, 1]$ is uncommitted, while $P(A) = [.76, .76]$ is consummate. Some have argued that indeterminacy captures "pre-systematic" notions of belief and disbelief [Sha76], [Lev80a].² Since $0 \leq \inf P(A) + \inf P(\sim A) \leq 1$, the agent can assign zero belief to a proposition even though he is not certain that it is false. Indeterminacy is useful to the subjectivist when eliciting bounds on probabilities (especially from equivocating experts), and to the empiricist for expressing the Neyman-Pearson confidence results of population sampling.

Intervalism is also natural in detachment. When

$\mathcal{K}(A, Q) = 1 - \delta$, then $\mathcal{K}(A, P \& Q) = [1 - \epsilon - \delta, 1]$.

If probabilities are based on direct inference from the class A , the probability of " $Px \& Qx$ " for some $x \in A$ would be an interval, despite having started with probabilities that were points (see [TSI83], [Che83], and [Nil84]).

Many advocates of interval belief measures in A.I. link their arguments to Shafer's interpretation [Sha76] of Dempster's inference system [Dem68]. Shafer's theory is claimed to provide a valuable representation of intervals (via mass functions), and a simple, consistent approach to resolving apparent disputes when combining evidence (via Dempster's rule). These claims are evaluated elsewhere [Kyb85], [Lev80a], [Zad79]. Shafer's theory is not unique in its ability to cope with disagreeing evidence; indeed, a system of belief would be impoverished if it made no provisions for disagreement (see Levi's remarks [Lev80a]; also, there are indeterminate systems due to Levi, Smith, Schick, Good, and Kyburg). Further, Dempster's rule for combining evidence is relatively

92-13702



presumptuous as a form of conditionalization [Dem68], [Lev80b], [Ky85].

Putting aside the prospects for Dempster's rule, we are left with these indeterminate probabilities, and with an ensuing decision problem. Barnett [Bar81] and Lowrance [Low82] have both suggested that a research goal should be a fully developed decision theory based on interval measures.

Luce and Raiffa call decision problems with indeterminate probabilities "partial ignorance" problems, and earlier work on partial ignorance is discussed in [Lou85]. Wesley, Lowrance, and Garvey [WLG84] offer a candidate theory that is for use with Shaferian beliefs and that ignores risk; it has been discussed elsewhere [Lou84].

II. Estimation and Decision.

With interval probabilities or interval utilities, expected utilities are intervals. If interval probabilities are narrow (or otherwise fortuitous) there is no problem: expected utility intervals can be ordered in the natural way (see below), and the best act identified. In a 1:1 lottery that depends on the outcome of a coin toss, if $P(\text{heads})$ is $[.7, .8]$ the decision should be clear via the obvious ordering; if it is $[.3, .8]$, the decision may not be clear. The decision may also not be clear if the interval is narrow, but unfortuitous, e.g., $[\cdot 49, \cdot 52]$.

If the maximization of expected utility (MEU) is the sole solution criterion, there may be no defensible ordering of the utility intervals that identifies a best act. Of course, MEU with point-probabilities can be ambiguous too. This latter ambiguity is often tolerated: if two acts have the exact same expected utility, the sameness of utility is supposed to reflect indifference. But ambiguity with interval probabilities may not be tolerable because intervals often model ignorance, not indifference. It is not the case that the two acts *couldn't* be ordered in a relevant and accountable way. Rather, not enough is known to order them.

There are two problems here. First, there is the estimation problem: what should be the degree of certainty attributed to a proposition? Second, there is

the decision problem: which act should be chosen among available acts, when the agent is not indifferent about them all? In the estimation problem, error is avoided by using intervals. In the decision problem, ambiguity is avoided by eschewing intervals. In order to solve both problems simultaneously, there must be some compromise.

III. Secondary Criterion Solutions.

Let Π be the largest set of probability distributions satisfying all of the interval constraints. Calculating expected utilities in the usual way, for act a_λ , in the presence of uncertainties E_i :

$$u_k(a_\lambda) = \sum \{P_k(E_i | a_\lambda) u(E_i | a_\lambda)\} \\ \forall E_i$$

$$U(a_\lambda) = \{u_k(a_\lambda) : P_k \in \Pi\}$$

$$u(a_\lambda) = [\inf U(a_\lambda), \sup U(a_\lambda)].$$

The natural way to (partial-) order acts with indeterminate utilities is by dominance: $a_1 > a_2$ iff $\inf U(a_1) > \sup U(a_2)$. If there is a unique maximal element in the order, a^* , then the decision problem is solved. The probabilities, though individually indeterminate, are nevertheless collectively *decisive*. But in general, there will be some set of maxima, $\{a_i\}$.

Some authors ([Hur51], [Goo83], [Fis65], [Lev80b]) suggest that a^* can be identified in the maximal set by one of the so-called weaker methods: maximin, min-regret, or lexi-min methods. These are the methods recommended for decision problems under uncertainty, and their common character that is crucial here is that they make no use of probability judgement. Presumably the probability information has been milked for all it is worth, under the primary method of MEU, and secondary methods will finish the job of identifying a^* . Unless there is an unforeseen equality of point-valued utilities, the weaker method guarantees identifying a unique act. The weaker method *could* have been applied in the first place but for its admitted weakness. It is considered weak precisely because it ignores probability judgement. It is employed secondarily precisely because it presupposes that probability judgement will be of no further use, which is exactly the case among the

Availability Codes	
Dist	Avail and/or Special
A-1	

maximal set after the application of MEU.

For programmed systems, there is still the problem of choosing one from among the various secondary criteria. Clearly there are situations in which maximin is inappropriate, and similarly for min-regret, for optimism-pessimism, etc. Supervaluations would be cautious, but impotent; taking the most popular mandate among various criteria would be ad hoc. One could attempt to discriminate those situations in which one method applies and others do not, but no such attempt has been successful.

IV. A Different Proposal.

Here is an entirely different way of solving partial ignorance problems. If MEU with the given probability intervals is indecisive, MFU can be retained and the probability intervals refined. In the Bayesian tradition, refinement of intervals is done subjectively, with no additional empirical information. In the Neyman-Pearson tradition, refinement is done objectively and requires additional empirical information. In either case, refinement further determines probabilities. An automatic reasoning system may be required to be objective, may not have recourse to additional information, and may require the preservation of indeterminacy. Fortunately, it's possible to refine intervals objectively, with no additional empirical information, and without losing the indeterminacy of probabilities. This latter possibility is presented more carefully in [1985].

Let credal state be described not by one set of feasible distributions, Π , but by a sequence of sets, $\langle \Pi_i \rangle$. Each Π_i is based on a body of knowledge formed with some quantifiable attitude toward error (so there is a companion sequence, $\langle K_i \rangle$, where each body of knowledge, K_i , has an integer index and a real error). Successive K 's are more informative, but predecessors are less prone to error. Each indeterminate expected utility calculation is done with respect to one element in the Π -sequence, but the whole Π -sequence constitutes the credal state. Therefore, different indeterminate probabilities, with different maximal sets, can be consulted for the

purposes of decision without changing the indeterminacy of the credal state.

This representation finesses the question of how narrow intervals ought to be. Imagine the expert who first reports the interval is $[.3, .7]$, but can be coaxed into reporting to the more useful $[.35, .65]$. Which interval gets represented? In this proposal, both should be represented. Intervals should be as narrow as permitted given the magnitude of error (α) associated with the body of knowledge on which the intervals are based, $K_{(\alpha)}$. They will be $[0, 1]$ in Π_0 . They may be degenerate in the very late Π 's. And they should be *variously* narrow (though not necessarily nested) in between.

In practice, this proposal requires additional represented information, or additional inference rules and epistemological assumptions. It may be possible simply to assert and to represent both sequences, $\langle \Pi_i \rangle$ and $\langle K_i \rangle$. But more likely, Π 's will have to be generated from K 's, and successive K 's from some initial base, K_{init} . A combination of the two methods, generation and assertion, is convenient.

Generating Π 's from K 's requires the adoption of some theory of probability. It could be as simple as taking statements in K to be constraints on distributions, or conditionalizing some prior on the contents of K , or it could be some theory of frequency-based or chance-based direct inference.

Generating $\text{succ}(K)$ amounts to making additional assumptions. It could be done in a number of ways: one possibility is to use an acceptance rule (see also below, on "higher-order" probabilities). Such a rule would describe when a statement is acceptable and would thus determine to which K 's it belongs. If the rule is based on probabilities relative to K_{init} , for instance, then A belongs to all those successor K 's, K_j , such that $1 - P(A | K_{init})$ is less than the error associated with K_j . A different probabilistic rule would take $\text{succ}(K)$ to be $K \cup \{A\}$, where A is the next most probable statement relative to K , of statements of some special form. Note that with these rules, the K -sequence is nested. Acceptance rules in the literature are more elaborate. See [Kyb70] for additional

acceptance rules and their evaluation.

Decision-making amounts to exploring the Π -sequence in best-first order until either (1) the maximal set under Π is a singleton, in which case the problem is solved; or (2) Π is a singleton, which leaves the standard decision problem under risk; or (3) the error associated with Π is intolerably large for this decision problem, in which case MFCU with no acceptable set of assumptions can legislate unambiguously.

The reasonableness of this proposal depends on whether there is independent reason to use intervals of a particular width. It may be that epistemological considerations require that certain intervals be used: e.g., the narrowest intervals at .95 confidence [Ky85]. But if not, if confidence levels .94 and .96 are also useable, then decision analysis might as well proceed with intervals that are decisive, rather than with those that are indecisive. There is no reason to avoid tolerable error if doing so results in uninformative analysis.³ If the MFCU calculation is not satisfactory under the assumptions held, it could be that the agent has not assumed enough. The analysis should then be founded on an augmented set of assumptions.

Conversely, there is no reason to invite error in the analysis if the analysis is already sufficiently informative. So of the many Π 's that are decisive, the one that is least prone to error has epistemic priority. The augmented set of assumptions should be the next-least in order of presumptiveness. No more assumptions should be made than are necessary for decision.

Consider the claim that rational commitment ceases with the restriction to the maximal set, or that the agent must sometimes suspend judgement when the set of maxima is not a singleton. Lopes, voicing a common intuition, quips that suspending judgement among choices with overlapping expected utility ranges is no more defensible than suspending judgement among choices with overlapping outcome ranges [Lop83]. Lopes' remark is forceful precisely because it points out the arbitrariness of interval width. Why invite error by using intervals narrower than [0, 1]? Because [0, 1] intervals are not satisfactorily

informative. But if the maximal set under narrower intervals is unsatisfactory, and the limit of tolerable potential error has not been reached, why not use still narrower intervals? One is already willing to forego certainty, and the amount of certainty one is willing to forego depends on the other desiderata, including decisiveness.

We still avoid error by using indeterminacy: we retain the early elements in the $\langle \Pi_i \rangle$ sequence, rather than settling immediately on the most specific element (or of some P , s.t. $P \in \Pi_i$ for all i). There may not be a most specific element in the sequence (this is explored in example C, below). And there may be genuine instances in which no substantiable set of assumptions legislates a unique decision or formulates a standard risk problem. In such cases, indeterminacy is required to indicate ignorance, or if either is possible, the need for more sampling, or for suspension of judgement.

For example, consider a probabilistic acceptance rule: statements are accepted in K_α when their probability relative to K_{init} exceeds $1 - \alpha$. For a decision problem where the maximum ratio of odds is w , it would be pointless to perform an MFCU analysis in some K_α where $\alpha \geq 1 - w$. If the lottery pays 20:1, $w = .95$. If all Π 's based on less error than .05 are indecisive, no decision is legislated (see [Ky85] and [Lou85] for discussion).

V. Examples and Contrasts.

We discuss the following decision problem.⁴ Upon finding a berry, the agent has to decide whether to eat it (a_1), or not to eat it (a_2). If it is eaten, it matters whether or not it was a good berry (G). If it is not eaten, it matters whether or not the agent later gets hungry (H). Let $u\langle a_1, G \rangle = 10$; $u\langle a_1, \sim G \rangle = -30$; $u\langle a_2, H \rangle = -10$; and $u\langle a_2, \sim H \rangle = 0$.

A. lower level confidence intervals.

Suppose the probability reports for G and for H are based on Clopper-Pearson intervals. Of 4 berries eaten, 4 were good. On 14 excursions of this kind, the agent got hungry (without eating) 3 times. At .99 confidence, $P(G) = [.35, 1]$ and $P(H) = [0, .55]$. So

$u(a_1) = [-16.8, 10]$; $u(a_2) = [-5.5, 0]$. The maximal set is $\{a_1, a_2\}$. But at the confidence level .75, $P(G) = [.75, 1]$ and $P(H) = [.15, .3]$. So $u(a_1) = [0, 10]$; $u(a_2) = [-3, -1.5]$. $a_1 \succ a_2$, a_1 is uniquely maximal. Note that if a_1 and a_2 had been ranked by utility midpoints at .99, $mp_1 = -3.4$; $mp_2 = -2.75$; one would have concluded contrarily that $a_2 \succ a_1$!

B. direct inference and probabilistic acceptance rule.

Suppose $\mathcal{U}(\text{berries, good}) = [.3, .8]$ and $\mathcal{U}(\text{excursions, get-hungry}) = [0, 1]$ and $\mathcal{U}(\text{soft berries, good}) = [.84, .88]$. Presumably this is accepted based on sampling with, say, at least .999 confidence. If $P(G)$ is based on the $[.3, .8]$ interval, both a_1 and a_2 are maximal. The decisive $[.84, .88]$ interval can't be used for $P(G)$ unless it is *accepted* that the berry is soft. Even if there is independent reason to believe $P(\text{this berry} \in \text{soft berries}) = .999$, the probability of G would be $[.3, .8]$. It's natural to consider the acceptance of "this berry \in soft berries". This allows direct inference: $P(G)$ must be $[.84, .88]$ if this is all that is known.⁵ Decision to do a_1 is based on dominance with the narrower interval.

C. convex Bayesian vs. Savage's Bayesian.

A Bayesian who considers all the distributions in a closed convex set can accept different constraints on this set at different levels of acceptance (cf. [Lev80b]). Typical constraints could be conditions (as in example B), or bounds on marginal probabilities (as in example A). Additional knowledge can lead to additional constraints, which can decrease membership in Π and so are more informative (though additional knowledge does not always lead to additional constraints: sometimes it can invalidate a constraint). Some constraints may not be as warranted as others, and their use introduces more possibility of error. If the set is indecisive, try the MEU analysis with the next set of constraints.

Savage would have the agent settle on the most specific set (if there is one), and eliminate the excess

indeterminacy of the preceding sets. If all the sets are nested (for all $i > j$, $\Pi_i \supseteq \Pi_j$), there is no difference between the decisions made by this convex Bayesian and by Savage's Bayesian.

But sets are not nested. The most obvious source of non-nesting is due to conditionalization.⁶ Suppose Π_1 is based on acceptance so stringent that probabilities are conditional only on A . Π_2 takes both A and B as conditions; B is acceptable as a condition at this level (perhaps B is treated by Jeffrey's rule in Π_1 ; it doesn't matter here). Then there's no reason for Π_2 to be a subset of Π_1 .

Let A entail $\sim H$: $P(G|A) = [.6, .8]$; $P(G|\sim A, B) = [.3, .4]$. Intuitively, A might be the conjunction "just ate & the berry looks good" while B might be "the lighting is misleading". Π_1 , with $P(G) = [.6, .8]$, indicates both a_1 (eating) and a_2 (not eating) as maximal. Π_2 mandates a_2 , with $P(G) = [.3, .4]$, which is not a sub-interval of $[.6, .8]$. Now suppose the next decision involving $P(G)$ is a 1:1 lottery. Π_1 mandates entering the lottery, and because Π_1 is decisive and epistemically prior, Π_2 is ignored. Savage would continue to use $P(G)$ from Π_2 , and would avoid the lottery.

So preservation of the "excess" indeterminacy is necessary despite temporary refinement for the purposes of the current decision.

D. Shaferian discounting.

It's tempting to consider Shafer's discounting parameter to generate successive Π 's.

The belief with mass $m(G) = .7$ and $m(\sim G) = .3$ is to be combined with a belief $m(G) = .6$; $m(\sim G) = .4$ based on a new, independent source. The latter's impact is to be discounted by some amount r . Let $\sim H$ be accepted. If $r < .23$, then $P(G) > .75$, and $a_1 = a^*$; otherwise $a_2 = a^*$. Note that for any value of r here, the resulting probability of G is determinate.

Are some values of r more cautious than others? If r is large, the informative impact of the second belief is lessened, and it is combined with caution. But a cautious attitude toward the new belief is not

necessarily a cautious attitude toward the possibility of error, unless the new belief is the only possible source of error. When conditions were not accepted in example C, it was because they were relatively uncertain, not because they were new. Here, it may be that the full weight of the new belief is required to avoid error. It would be erroneous, for instance, to ignore the new belief completely. The parameter r here is being used like Carnap's λ . There is no epistemic relation given between r and error, hence, no priority of one solution over the other.

Perhaps caution should be reflected by discounting both belief functions. This begs the question, in what proportion should they be discounted? If there are two parameters that can be varied, the Π 's generated will be only partially ordered.

It may be possible to use Shafer's formalism to generate the Π -sequence, but its use would require more argument.

VI. Epistemological Considerations.

A. *On Revisions of the Knowledge Base.*

A behavioral interpretation of probability suggests the identification of a^* as additional evidence about probability judgement. Whatever the means of a^* 's identification, there is a set, Ψ , of admissible probability distributions, according to each of which, a^* is the unique maximum by MEU means alone. Behaviorists hold that once a^* is identified, the agent's credal state contracts to the more precise Π , the intersection of Ψ and Π , at least as a description appropriate at the time of decision. Presumably, if there is no subsequent revision, the more precise description of past state continues to describe the current state. If this is right, then credal state depends on the decisions made. Faced with a different decision structure, a^* , hence Ψ , and finally credal state, might have been different.

Upon each decision, the agent must be consistent, in this behaviorally strong sense. Decisions always reveal credal state and always do so through

MEU.

There are enormous implications of this revelation-through-behavior stance for the management of knowledge bases. No matter how tentative the decision, and whatever its content or manner of selection, the knowledge base must represent only the distributions that are MEU-admissible for that decision. If only a single distribution is MEU-admissible, then that distribution specifies the new state of the program's belief. And this has been done with the addition of no relevant empirical knowledge! All that distinguishes the new state from the old is the actualization of one particular problem structure, among the many that could have been faced.

If the interpretation of probability is subjective as well as behavioral, the agent or reasoning system can spuriously return to the more permissive credal state, Π . But if this is to be a rule for revision, there seems no point in making the contraction. If it is not a rule, then there is still the onerous possibility of spurious change to some other credal state, and worse, the possibility of no change whatsoever after contraction.

Either course violates legitimate counterfactual intuitions pertaining to the past decision. Suppose the secondary method is always a tournament of coin-flipping. Upon the last toss of heads, a_2 is chosen, and Π' is obtained from Π by the deletion of all distributions that do not mandate a_2 . Thus, it no longer is the case that "had the toss been tails, a_1 would have been mandated," though we quite reasonably take such to have been the case.

Starr [Sta66] suggests a normative criterion for identifying the optimal act when Π is not a singleton. Suppose the distributions in Π can be parameterized by some θ . Suppose also that the set of parameter values Θ , corresponding to the Π distributions, is measured by an additive indifference "prior". So subsets of Π are also measured. Consider various acts. An act is mandated by each of its MEU-admissible distributions, which collectively form some subset $\pi \subseteq \Pi$. Starr's criterion chooses the act with the π that maximizes the measure (i.e., that has the greatest

number of feasible admissible distributions).

Starr's criterion is a prescription for decision, not for the adoption of a narrower credal state. Behaviorists would contract to π .

Whatever the behaviorist arguments, the revelation of credal state through decisions and MEU is unattractive in A.I. A system's probability estimates are based on objective analysis of samples, or on the opinions of experts, not on the future decision problems to be faced by the system.

B. On Higher-Order Probabilities.

Some Bayesians intuit the existence of "higher order" probabilities (e.g., [Goo83]). These would be probability distributions on probability distributions, formalized perhaps, like the indifference "prior" in Starr's criterion.

If one approves of and has access to such measures, then acceptance can be based on the measure. For instance, successive Π 's could be generated by eliminating the next-least probable members of the previous Π . This strategy leads to nested Π 's; all decisions would be those mandated by the distribution with the greatest higher-order-measure. It would not, in general, be the same as taking an expectation over the expected utility intervals, and ranking the resulting real-values:

$$u(a_\lambda) = \sum \{ \sum [P_k(E_i | a_\lambda) u(E_i | a_\lambda)] M(P_k) \} \\ \forall k \forall i, \\ \text{where } M \text{ is the higher-order measure.}$$

Perhaps the expectation is appropriate if there is such a measure. However, one should have misgivings about the identification of these measures.

There may be uncertainty about the higher-order measure, reflected in some still higher measure. This induces a hierarchy of measures. Presumably the height of the hierarchy is finite. There must be, at some high order, either a determinate measure, or else unmeasured indeterminacy. If the former, then one should be suspicious about the source of a determinate higher-order-measure: why is the probability of a distribution certain, but the distribution uncertain? The higher-level is not inherently more robust (note

that the order of the sums can be reversed). Just as a small error in a probability can change a decision, so can a small error in a higher-order probability change a decision.

If on the other hand there is unmeasured indeterminacy, the expected-expected utilities will be intervals. This is essentially no different from the interval expected utilities from indeterminate zero-order probabilities.

So acceptance can be conceptually related to higher order probabilities, but is not immediately subsumed or improved by them.

VII. Conclusion.

A.I. systems that use interval judgements must sometimes solve partial ignorance decision problems. There are now two approaches. Maximizing expected utility can be followed by maximin, or some other secondary criterion. Alternatively, additional assumptions can be made that change probabilities, temporarily, so that maximizing expected utility is sufficient. This paper has discussed how to implement the latter approach. Assumptions are accepted in an order that tries to avoid error, and they are accepted only temporarily, for the purposes of decision.

There is still the problem of choosing an acceptance rule, which iteratively generates the next-best assumption. This choice requires considerably more epistemological reflection.

VIII. Notes.

¹ By *indeterminacy*, we will mean indeterminacy broadly construed: potential indeterminacy, including the judgements $\text{Prob}(A) \in [0, 1]$ (complete) and $\text{Prob}(A) \in [1, 1]$ (degenerate), $\text{Prob}(A) \in [3, 7]$ (bounded), and $\{\text{Prob}(A) = .4 \text{ or } \text{Prob}(A) = .8\}$ (disjunctive).

² Some have charged that the specification of an interval requires two numbers rather than one; hence, it requires more information. That's silly. Given that some quantity p is in fact 0.67, it follows that p is in the interval $[0.34, 0.97]$. Furthermore, in a very natural canonical form, namely, the number of hyperplanar constraints required in the space of all probability distributions: the information (number of constraints) in interval reports of a particular probability is less than the information in point reports. Information measures are dependent on canonical form, hence can be misleading.

Intervals are chosen because they offer robust behavior. If practice shows that they are not robust enough, that endpoints matter critically, then future investigators can feel free to use a formalism with indeterminate upper and lower bounds, or with fuzzy sets. Surely one would not revert to point probabilities because they contain "less information."

³ Here, we've taken informativeness w.r.t. decision to be singularity of Π or singularity of the maximal set. Other interpretations of "informative" are possible (such as any restriction of the maximal set to decisions which cannot differ in outcome more than ϵ). These lead to different decision theories.

Also note that in [Lou85], the amount of tolerable error is addressed (see the discussion of D-meaningful corpora).

⁴ We call this the problem of Jerry's Bernies.

⁵ We've appealed to the epistemological conception of probability here. If explicit statement of chances is required, the example can be changed.

⁶ It's also possible to violate nesting when constraints are ordered jointly, and not all constraints are compatible. So if c_1, \dots, c_4 are constraints on Π 's, Π_1 may be delimited by $\{c_1\}$, and Π_2 by $\{c_2\}$, and Π_3 by $\{c_3\}$ before Π_3 by $\{c_1, c_2\}$. Π_5 may be delimited by $\{c_1, c_4\}$, where $\{c_1, c_2, c_4\}$ is overdetermining. If constraints are accepted (rather than knowledge that generates constraint), and acceptance is purely probabilistic, then this kind of situation requires acceptance levels at or below .5. With not purely probabilistic acceptance, this situation is more natural.

Note that non-nested Π 's would seem irrational via a Dutch Book argument, but the agent still posts consistent odds whenever he considers two or more lotteries simultaneously. It's only when he posts odds independently and they are subsequently collected that leads to inconsistency. Consult the Ellsberg paradox for intuitions here.

IX. References

- [Bar81] Barnett, J. "Computational Methods for a Mathematical Theory of Evidence." IJCAI 7.
- [Che83] Cheeseman, P. "A Method of Computing Generalized Bayesian Probability Values for Expert Systems." IJCAI 8.
- [Dem68] Dempster, A. "A Generalization of Bayesian Inference." *J. Roy. Stat. Soc. 2*.
- [Dil82] Dillard, R. "The Dempster-Shafer Theory Applied to Tactical Fusion in an Inference System." Fifth MIT/ONR Workshop.
- [Fis65] Fishburn, P. "Analysis of Decisions with Incomplete Knowledge of Probabilities." *Op. Res. 13*.
- [GLF81] Garvey, T., Lowrance, J., and Fischler, M. "An Inference Technique for Integrating Knowledge from Disparate Sources." IJCAI 7.
- [Gin84] Ginsberg, M. "Non-monotonic Reasoning using Dempster's Rule." AAAI.
- [Goo83] Good, I. *Good Thinking: The Foundations of Probability and Its Applications*. Minnesota.
- [GoM84] Good, I. and McMichael, A. "A Pragmatic Modification of Explicativity for the Acceptance of Hypotheses." *Phil. Sci. 51*.
- [Hur51] Hurwicz, L. "Some Specification Problems and Applications to Econometric Models." *Econometrica 19*.
- [Kyb70] Kyburg, H. "Conjunctivism," in Swain, M. ed. *Induction, Acceptance, and Rational Belief*. Dordrecht: Reidel.
- [Kyb74] Kyburg, H. *The Logical Foundations of Statistical Inference*. Dordrecht: Reidel.
- [Kyb85] Kyburg, H. *Science and Reason*. manuscript.
- [Lev80a] Levi, I. "Acceptance as a Basis for Induction," in Cohen and Hesse, eds., *Applications of Inductive Logic*. Oxford.
- [Lev80b] Levi, I. *The Enterprise of Knowledge*. MIT, Cambridge.
- [Lop84] Lopes, L. "Normative Theories of Rationality: Occam's Razor, Procrustes' Bed?" (response to Kyburg). *Behav. Brain Sci. 6*.
- [Lou84] Loui, R. "Primer for the Inexact Reasoning Session." memorandum.
- [Lou85] Loui, R. "Decisions with Indeterminate Probabilities." submitted.
- [Low82] Lowrance, J. "Dependency-Graph Models of Evidential Support." Ph.D. Thesis, U. Mass. Amherst. COINS TR.
- [LuS84] Lu, S. and Stephanou, H. "A Set-theoretic Framework for the Processing of Uncertain Knowledge." AAAI.
- [Nil84] Nilsson, N. "Probabilistic Logic" SRI Tech. Note 321.
- [Qui83] Quinlan, J. "Consistency and Plausible Reasoning." IJCAI 8.
- [Sha76] Shafer, G. *A Mathematical Theory of Evidence*. Princeton.
- [Sta66] Starr, M. "A Discussion of Some Normative Criteria for Decision-Making under Uncertainty." *Industrial Management Review 8*.
- [Str84] Strat, T. "Continuous Belief Functions for Evidential Reasoning." AAAI.
- [TSD83] Tong, R., Shapiro, D., Dean, J., and McCune, B. "A Comparison of Uncertainty Calculi in an Expert System for Information Retrieval." IJCAI 8.
- [WeH82] Wesley, L. and Hanson, A. "The Use of an Evidential Based Model for Representing Knowledge and Reasoning about Images in the Vision System." *IEEE*.
- [Wes83] Wesley, L. "Reasoning about Control: The Investigation of an Evidential Approach." IJCAI 8.
- [WLG84] Wesley, L., Lowrance, J., and Garvey, T. "Reasoning about Control: An Evidential Approach." SRI Tech. Note 324.
- [Zad79] Zadeh, L. "On the Validity of Dempster's Rule." U.C. Berkeley/ERL M79/24.

Interval-Based Decisions for Reasoning Systems

Ronald P. Loui*

Departments of Computer Science and Philosophy
University of Rochester
Rochester, NY 14627

*originally presented in conjunction with
Jerome A. Feldman and Henry G. E. Kyburg, Jr.

This essay looks at decision-making with interval-valued probability measures. Existing decision methods have either supplemented expected utility methods with additional criteria of optimality, or have attempted to supplement the interval-valued measures. We advocate a new approach, which makes the following questions moot: 1. which additional criteria to use, and 2. how wide intervals should be. In order to implement the approach, we need more epistemological information. Such information can be generated by a *rule of acceptance* with a parameter that allows various attitudes toward error, or can simply be declared.

In sketch, the argument is: 1. probability intervals are useful and natural in A.I. systems; 2. wide intervals avoid error, but are useless in some risk-sensitive decision-making; 3. one may obtain narrower intervals if one is less cautious; 4. if bodies of knowledge can be ordered by their caution, one should perform the decision analysis with the acceptable body of knowledge that is the most cautious, of those that are useful. The resulting behavior differs from that of a behavioral probabilist (a Bayesian) because in the proposal, 5. intervals based on successive bodies of knowledge are not always nested; 6. if the agent uses a probability for a particular decision, she need not commit to that probability for credence or future decision; and 7. there may be no acceptable body of knowledge that is useful; hence, sometimes no decision is mandated.

1. Interval Measures.

Interval measures for credal probabilities have taken A.I. by storm. Researchers selecting formalisms for quantifying belief have recently recognized the virtues of indeterminacy in probability judgement ([1], [4], [6], [7], [18], [19], [21], [24], [26], [27], etc.).

Intervals allow varying degrees of commitment in probability assertion. At the extremes, $Prob(A) = [0, 1]$ is uncommitted, while $Prob(A) = [.76, .76]$ is consummate. Some have argued that indeterminacy captures "pre-systematic" notions of belief and disbelief ([14], [22]).¹ Since $0 \leq \inf Prob(A) + \inf Prob(\sim A) \leq 1$, the agent can assign no belief to a proposition even though she is not certain that it is false. Indeterminacy is useful to the subjectivist when eliciting bounds on probabilities (especially from equivocating experts), and to those who take interval estimation seriously.

Intervalism is also natural in detachment (which leads to so-called probabilistic logics). When the relative frequency of P 's among A 's, $\% (A, P)$, is $1 - \epsilon$, and the relative frequency of Q 's among A 's, $\% (A, Q)$, is $1 - \delta$, then with deductive certainty, the relative frequency of $(P \& Q)$'s among A 's, $\% (A, P \& Q)$, is $[1 - \epsilon - \delta, 1]$. If probabilities are inferred directly from the class A , the probability of " $P \& Q$ " for some $x \in A$ will be an interval, despite having started with probabilities that were points (see [2], [20], and [25]).

Many advocates of interval belief measures in A.I. link their arguments to Shafer's interpretation [22] of Dempster's inference system [3]. Shafer's theory is claimed to provide a valuable representation of intervals (via mass functions), and a simple, consistent approach to resolving disputes when combining evidence (via Dempster's rule). These claims are evaluated elsewhere ([13], [15], [29]). Shafer's theory is not unique in its ability to cope with disagreeing evidence; indeed, a system of belief would be impoverished if it made no provisions for disagreement (see Levi's remarks [14]; also, there are indeterminate systems due to Levi, Smith, Schick, Good, and Kyburg). Further, Dempster's rule for combining evidence is relatively presumptuous as a form of conditionalization.

Putting aside the prospects for Dempster's rule, we are left with these indeterminate probabilities, and with an ensuing decision problem. Barnett [1] and Lowrance [18] have both suggested that a research goal should be a fully developed decision theory based on interval measures. At this conference, and at the IJCAI that followed, the possibility of an effective, interval-based decision theory was a central point of dispute.

Decision-theorists have of course studied the problem. Luce and Raiffa call decision problems with indeterminate probabilities "partial ignorance" problems, and earlier work on partial ignorance is discussed in [17].

2. Estimation and Decision.

With interval probabilities, expected utilities are intervals. Various ways of calculating utilities with interval probabilities have been proposed. Most authors retain the

standard definition of expected utility for real-valued distributions, then focus on the set of real-valued probability distributions satisfying the interval constraints. This leads to a set of expected utilities, each generated by a real-valued distribution.

Let Π be the largest set of probability distributions satisfying all of the interval constraints. Let $\{a_\lambda\}$ be acts under consideration and $\{E_i\}$ be exclusive and exhaustive events. u is a utility measure over tuples of events and acts.

$$u_K(a_\lambda) = \sum_{\forall E_i} \{Prob_K(E_i | a_\lambda) u(E_i, a_\lambda)\}$$

is the usual real-valued expected utility for the real-valued measure $Prob_K$.

$$\underline{U}(a_\lambda) = \{u_K(a_\lambda) : Prob_K \in \Pi\}$$

is the set of these, for probability measures belonging to Π .

$$\underline{u}(a_\lambda) = [\inf \underline{U}(a_\lambda), \sup \underline{U}(a_\lambda)]$$

is the expected utility interval, determined by bounding the set of real-valued expected utilities.

The natural way to order acts with indeterminate utilities is by dominance: $a_1 \succ a_2$ iff $\inf \underline{U}(a_1) > \sup \underline{U}(a_2)$. If the maximal element in the order is unique, then call it a^* , and the decision problem is solved. The probability, though indeterminate, is decisive. In general, there will be some set of maxima, $\{a_j\}$, which we'll call the maximal set.

If interval probabilities are narrow, or if the bounds fall in the appropriate places, there is no problem: expected utility intervals can be ordered in the natural way, and the best act identified. In a fair lottery based on the toss of a biased coin, where bias is estimated indeterminately, if $Prob(heads)$ is $[.7, .8]$ the decision should be clear. If it is $[.3, .8]$, the decision is unclear. The decision may also be unclear if the interval is narrow, but unfavourable, e.g., $[.49, .52]$.

If the maximization of expected utility (MEU) is the sole solution criterion, there may be no defensible ordering of the utility intervals that identifies a best act. Of course, MEU with point probabilities can be ambiguous too. This latter ambiguity is often tolerated; if two acts have the exact same expected utility, the sameness of utility is supposed to reflect indifference. But ambiguity with interval probabilities may not be tolerable because interval probabilities model ignorance, not indifference. It is not the case that the two acts *couldn't* be ordered in a relevant and accountable way. Rather, not enough is known to order them.

There are two problems here. First, there is the estimation problem: what should be

the rational estimate of the probability of an event? Second, there is the decision problem: which act should be chosen among available acts, when the agent is not indifferent among them all? In the estimation problem, intervals avoid error. In the decision problem, real-valued probabilities are needed to avoid ambiguity. Solving both problems simultaneously requires some compromise.

3. Secondary Criterion Solutions.

Some authors ([5], [8], [9], [15]) suggest that a^* , the optimal act, can be identified in the maximal set by one of the so-called weaker methods: min-regret, optimism-pessimism, or lexi-min methods. These are the methods recommended for decision problems under uncertainty, and their common character that is crucial here is that they make no use of probability judgement. Presumably the probability information has been used for all it is worth, under the primary method of MEU, and secondary methods will finish the job of identifying a^* . The weaker method guarantees identifying a unique act, or a set of acts between any of which the agent should be indifferent. The weak rule *could* have been applied before the set of acts was pruned to the maximal set. But it is weak; it ignores probability judgement. It is employed secondarily only because probability judgement is of no use among the maximal set.

For agents that repeatedly use one rule, including all programmed systems, there is still the problem of choosing one from among the various secondary criteria. If optimism-pessimism is used, what should be the tradeoff ratio, alpha? No single rule seems appropriate for use on every maximal set. Acting only when there is unanimous opinion would be cautious, but impotent; taking the most popular mandate among various criteria would be ad hoc. One could attempt to discriminate those situations in which one rule applies and others do not, but no such attempt has been successful.

4. A Different Proposal.

There is an entirely different way of solving partial ignorance problems. If MEU with the given probability intervals is indecisive, MEU can be retained and the probability intervals refined. Some hold that intervals can be refined by the agent's introspection, subjectively, with no additional empirical information. Others hold that refinement of intervals can only be done objectively via further empirical investigation. A.I. belief systems may be required to be objective, may not have recourse to additional information, and may require the preservation of indeterminacy. Fortunately, in our proposal, it's possible to refine intervals objectively, with no additional empirical information, and without losing the indeterminacy of probabilities.

Let credal state be described not by one set of feasible distributions, Π , but by a sequence of sets, $\langle \Pi_0, \dots \rangle$. Each Π_i is based on a body of knowledge, K_i (so there is a companion sequence, $\langle K_0, \dots \rangle$). Each pair $\langle \Pi_i, K_i \rangle$ is adopted by the agent with

some boldness, caution, or attitude toward error. Successive K_i 's are less cautious, but usually more informative. Caution is at least ordinal, and may even be representable by a real-valued measure (but not by a probability measure). Indeterminate expected utility calculations can be done with respect to any one element in the Π -sequence, but the whole Π -sequence constitutes the credal state. Therefore, different indeterminate probabilities, with different maximal sets, can be consulted for the purposes of decision without changing the credal state in any way.

This representation finesses the question of how narrow intervals ought to be. Imagine the expert who first reports the interval for a probability to be $[.3, .7]$, but who can be coaxed into reporting the more useful $[.35, .65]$. Which interval gets represented? In this proposal, both should be represented. Intervals should be as narrow as permitted given the boldness associated with the body of knowledge on which the intervals are based. They will be $[0, 1]$ in Π_0 . They may be degenerate, real-values in the very late Π_i 's. And they should be variously narrow (though not necessarily nested) in between.

In practice, this proposal requires additional represented information, or additional inference rules and epistemological assumptions. It may be possible simply to assert and to represent both sequences, $\langle \Pi_0, \dots \rangle$ and $\langle K_0, \dots \rangle$. But more likely, Π_i 's will have to be generated from K_i 's, and successive K_i 's from some initial base, K_0 . A combination of the two methods, generation and assertion, is probably best for implementations.

Generating Π_i 's from K_i 's requires the adoption of some theory of probability. Statements in K_i could simply be taken as constraints on distributions. The contents of K_i could be used to conditionalize some prior. One could use a full theory of frequency-based or chance-based direct inference. Or one could take the contents of K_i to be sampling information. If there is already a sequence of different bodies of knowledge, there will be a diversity among sets of probability distributions. Another possibility is to consider the use of successively bolder evidence combination rules to generate successive Π_i 's, for fixed K . Dempster's rule gives bounds narrower than those yielded by a classical indeterminate Bayesian analysis, so it can be considered a bold way to generate probabilities. Also, for fixed sampling information, nested Π_i 's can be generated by adjusting the confidence level for Clopper-Pearson estimates, or for estimates from some other confidence system that generates connected intervals.

Generating $\text{succ}(K_i)$, successive K 's, amounts to making additional assumptions. It too could be done in a number of ways. One possibility is to use an acceptance rule. Such a rule would describe when a statement is acceptable and would thus determine to which K_i 's it belongs. If the rule is based on probabilities relative to K_0 , for instance, then a statement A belongs to all those successors, K_i , such that $1 - \text{Prob}(A | K_0)$ is less than the threshold associated with K_i . A different probabilistic rule would take

$\text{succ}(K)$ to be $K \cup \{A\}$, where A maximizes $\text{Prob}(A | K_0)$, among a not in K , and a of some special form. Acceptance rules defended in the literature are more elaborate (see [10]). The K -sequence need not be nested, and the Π -sequence almost certainly won't be. Another possibility is to base successive K_i 's on increasingly large (i.e., bold) sets of default rules.

Decision-making amounts to exploring the Π -sequence in best-first order until, at some position in the sequence i , either (1) the maximal set of acts generated using probabilities in Π_i is a singleton, in which case the problem is solved; or (2) Π_i is a singleton, which leaves the standard decision problem under risk; or (3) the boldness associated with Π_i is intolerably large for this decision problem, in which case MEU with no acceptable set of assumptions can legislate unambiguously.

The reasonableness of this proposal depends on whether there is independent justification for intervals of a particular width. It may be that epistemological considerations require that certain intervals be used: e.g., the narrowest intervals at .95 confidence. But if not, if confidence levels .94 and .96 are also useable, then decision analysis might as well proceed with intervals that are decisive (that lead to singleton maximal sets), rather than with those that are indecisive. There is no reason to avoid tolerable error if doing so results in uninformative analysis.² If the MEU calculation is not satisfactory under the assumptions held, it could be that the agent has not assumed enough. The analysis should then be founded on an augmented set of assumptions.

Conversely, there is no reason to invite error in the analysis if the analysis is already sufficiently informative. So of the many Π_i 's that are decisive, the one that is least prone to error has epistemic priority. The augmented set of assumptions should be the next-least in order of presumptiveness. No more assumptions should be made than are necessary for decision.

Consider the claim that rational commitment ceases with the restriction to the maximal set, or that the agent must sometimes suspend judgement when the set of maxima is not a singleton. Lopes, voicing a common intuition, quips that suspending judgement among choices with overlapping expected utility ranges is no more defensible than suspending judgement among choices with overlapping outcome ranges [16]. Lopes' remark is forceful precisely because it points out the arbitrariness of interval width. Why invite error by using intervals narrower than $[0, 1]$? Because $[0, 1]$ intervals are not satisfactorily informative. But if the maximal set under narrower intervals is unsatisfactory, and the limit of tolerable potential error has not been reached, why not use still narrower intervals? One is already willing to forego certainty, and the amount of certainty one is willing to forego depends on the other desiderata, including decisiveness.

We still avoid error by using indeterminacy; we retain the early elements in the $\langle \Pi_0, \dots \rangle$ sequence, rather than settling immediately on the most specific element (or on some P_i , s.t. $P_i \in \Pi_i$ for all i). There may not be a most specific element in the sequence (this is explored in examples 5.2 and 5.3, below). And there may be genuine instances in which no substantiable set of assumptions legislates a unique decision or formulates a standard risk problem. In such cases, indeterminacy is required to indicate ignorance, or if either is possible, the need for more sampling, or for suspension of judgement.

For an example of a case where no decision is mandated, consider $\langle K_i, \Pi_i \rangle$ pairs constructed by a probabilistic acceptance rule: statements are accepted in K_i when their probability relative to K_0 exceeds $1 - f(i)$, for some f mapping integers to reals. The maximum ratio of odds in the decision problem is w . It would be pointless to perform an MEU analysis in some K_i where $f(i) \geq 1 - w$. If the lottery pays 20:1, $w = .95$. If all Π_i 's based on thresholds greater than .95 are indecisive, no decision is legislated (see [17] for discussion).

5. Examples and Contrasts.

We discuss the following decision problem. Upon finding a berry, the agent has to decide whether to eat it (a_1), or not to eat it (a_2). If it is eaten, it matters whether or not it was a good berry (G). If it is not eaten, it matters whether or not the agent later gets hungry (H). Let $u\langle a_1, G \rangle = 10$; $u\langle a_1, \sim G \rangle = -30$; $u\langle a_2, H \rangle = -10$; and $u\langle a_2, \sim H \rangle = 0$.

5.1. Confidence Levels.

Suppose the probability reports for G and for H are based on Clopper-Pearson intervals. Suppose also that normative theory requires preference according to the given utilities weighed by Clopper-Pearson estimates (the most natural way is to suppose that rationality requires that an agent's subjective probabilities fall within these interval estimates). Of 4 berries eaten, 4 were good. On 14 excursions of this kind, the agent got hungry (without eating) 3 times. At .99 confidence, $Prob(G) = [.35, 1]$ and $Prob(H) = [0, .55]$. So $u(a_1) = [-16.8, 10]$; $u(a_2) = [-5.5, 0]$. The maximal set is $\{a_1, a_2\}$. But at the confidence level .75, $Prob(G) = [.75, 1]$ and $Prob(H) = [.15, .3]$. So $u(a_1) = [0, 10]$; $u(a_2) = [-3, -1.5]$. $a_1 \succ a_2$. a_1 is uniquely maximal. Note that if a_1 and a_2 had been ranked by utility interval midpoints at .99, a popular secondary method, then $mp_1 = -3.4$; $mp_2 = -2.75$. One would have concluded contrarily that $a_2 \succ a_1$.

5.2. Direct Inference and Acceptance.

Suppose $\%(berries, good) = [.3, .8]$ and $\%(excursions, get-hungry) = [0, 1]$ and $\%(soft berries, good) = [.84, .88]$. These statements are accepted via the high probability rule at a level, say .999, relative to evidence (see [11]). Most theories of direct inference would base $Prob(G)$ on the $[.3, .8]$ interval, in which case both a_1 and a_2 are maximal. The decisive $[.84, .88]$ interval can't be used for $Prob(G)$ unless it is accepted that the berry is soft. Even if there is independent reason to believe $Prob(\text{this berry} \in \text{soft berries}) = .998$, the probability of G would be $[.3, .8]$. It's natural to consider a lower acceptance level, .998, at which "this berry \in soft berries" can be accepted. This allows direct inference from the narrower class: $Prob(G)$ must be $[.84, .88]$ if this is all that is known. The optimal decision is unambiguously a_1 with the narrower interval. Note that the Π just satisfying the constraints at .998 doesn't nest in the Π just satisfying the constraints at .999.

5.3. Convex Bayesian vs. Savage's Bayesian.

A Bayesian who considers all the distributions in a closed convex set can accept different constraints on this set with different degrees of boldness (e.g. Levi [15], for whom boldness, or caution, is a technical term). Typical constraints could be bounds on marginal probabilities (as in example 5.1), or conditions (as in example 5.2). Additional knowledge can lead to additional constraints, which can decrease membership in Π (though additional knowledge does not always lead to additional constraints; sometimes it can invalidate a constraint, especially if constraints are inferred defeasibly). Some constraints may not be as warranted as others, and their use introduces more possibility of error. If Π is indecisive, try the MEU analysis with the next set of constraints on Π , in order of increasing boldness (decreasing caution).

Savage would have the agent settle on the most specific set of constraints (if there is one), and ignore the preceding sets. If all the constraint sets are nested (for all $i > j$, $\Pi_i \supseteq \Pi_j$), there is no difference between the decisions made by this convex Bayesian and by Savage's Bayesian (though the convex Bayesian can sometimes suspend judgement while Savage's never can).

But successive Π_i 's are not nested. The most obvious source of non-nesting is due to conditionalization.³ Suppose Π_1 is based on epistemic policies so cautious that only A can be taken as a condition. Suppose Π_2 takes both A and B as conditions; B can be used as a condition at this level of caution, but not at the previous level. Then there's no reason for Π_2 to be a subset of Π_1 .

Let A entail $\sim H$ to simplify the decision problem. $Prob(G | A) = [.6, .8]$; $Prob(G | A, B) = [.3, .4]$. A might be the conjunction "just ate & the berry looks good" while B might be "the lighting is misleading". Π_1 , with $Prob(G) = [.6, .8]$, indicates both a_1 (eating)

and a_2 (not eating) as maximal. Π_2 mandates a_2 , with $\text{Prob}(G) = [.3, .4]$. Again the Π_1 's aren't nested since $[.3, .4]$ is not a sub-interval of $[.6, .8]$. Now suppose the next decision involving $\text{Prob}(G)$ is a 1:1 lottery. Π_1 mandates entering the lottery, and because Π_1 is decisive and epistemically prior, Π_2 is ignored. Savage would continue to use $\text{Prob}(G)$ from Π_2 , and would avoid the lottery.

So preservation of the indeterminacy represented by the earlier Π_1 's is necessary despite the temporary refinements used to make particular decisions.

5.4. Shaferian Discounting.

It's tempting to consider Shafer's discounting parameter to generate successive Π_1 's. Once again, suppose that normative theory requires preference according to the given utilities weighed by Shaferian belief.

The belief with mass $m(G) = .7$ and $m(\sim G) = .3$ is to be combined with a belief $m(G) = .6$; $m(\sim G) = .4$ based on a new, independent source. The latter's impact is to be discounted by some amount r . Let $\sim H$ be given, again to simplify the decision problem. If $r < .23$, then $P^*(G) = \text{Bel}(G) > .75$, and $a_1 = a^*$; otherwise $a_2 = a^*$. Note that for any value of r here, the resulting probability of G is determinate.

Are some values of r more cautious than others? If r is large, the informative impact of the second belief is lessened, and it is combined with caution. But a cautious attitude toward the new belief is not necessarily a cautious attitude toward the possibility of error, unless the new belief is the only possible source of error. When conditions were not accepted in example C, it was because they were relatively uncertain, not because they were new. Here, it may be that the full weight of the new belief is required to avoid error. It would be erroneous, for instance, to ignore the new belief completely. The parameter r here is being used like Carnap's λ . There is no epistemic relation given between r and error or caution, hence, no priority of one solution over the other.

Perhaps caution should be reflected by discounting both belief functions. This begs the question, in what proportion should they be discounted? If there are two parameters that can be varied, the Π_1 's generated will be only partially ordered. There is a natural extension of our proposal to partially ordered Π -sequences: look for agreement among maximal Π_1 's; but it doesn't seem very powerful.

It may be possible to use Shafer's formalism to generate the Π -sequence, but its use would require more argument. If caution can't be varied in this formalism, then users of Dempster's rule are in a curious position: they can get strictly narrower intervals than the convex Bayesian, but they can't choose how narrow the intervals should be. Shafer himself has formalized decision-making in a different way ([23]; see also [28]).

6. Epistemological Considerations.

6.1. On Revisions of the Knowledge Base.

A behavioral interpretation of probability suggests that identifying the best act in a decision problem should be taken as additional evidence about probability judgement. Whatever the means of identifying the optimal act, a^* , there is a set, Ψ , of admissible probability distributions according to each of which a^* is the unique maximum by MEU means alone. Behaviorists hold that once a^* is identified, the agent's credal state contracts to the more precise Π^* , the intersection of Ψ and Π , at least as a description appropriate at the time of decision. Presumably, if there is no subsequent revision, the more precise description of past state continues to describe the current state. If this is right, then the current credal state depends on the decisions made in the past. Faced with a different decision structure, a^* , Ψ , and credal state, might all have been different.

Every decision or preference of the agent must be consistent, in this behaviorally strong sense. Decisions always reveal credal state and always do so through MEU.

There are enormous implications of this revelation-through-behavior stance for the management of knowledge bases. No matter how tentative the decision, and whatever the stakes it involves or its manner of selection, the knowledge base must represent only the distributions that are MEU-admissible for that decision. If only a single distribution is MEU-admissible, then that distribution specifies the new state of the program's belief. And this has been done with the addition of no relevant empirical knowledge! All that distinguishes the new state from the old is the actualization of one particular problem structure, among the many that could have been faced.

If the interpretation of probability is subjective, the agent or reasoning system can spuriously return to the more permissive credal state, Π . But if this is to be a rule for revision, there seems no point in making the contraction. If it is not a rule, then there is still the onerous possibility of spurious change to some other credal state, and worse, the possibility of no change whatsoever after contraction.

Either course violates legitimate counterfactual intuitions pertaining to the past decision. Suppose the secondary method is always a tournament of coin-flipping. Upon the last toss of heads, a_2 is chosen, and Π^* is obtained from Π by the deletion of all distributions that do not mandate a_2 . Thus, it no longer is the case that "had the toss been tails, a_1 would have been mandated," though we quite reasonably take such to have been the case.

Philosophers have studied other defects of the behavioral stance and the Dutch Book arguments on which they depend (See [12] for a partial bibliography). And some have

argued for non-behavioral ways of looking at decision-making (see [15]). Whatever the philosophical point, the behavioral position seems unattractive for A.I.. A system's probability statements should be based on objective analysis of samples, or on the opinions of experts, not on the fact that the system made particular decisions in the past.

6.2. On Higher-Order Probabilities.

Some Bayesians intuit the existence of "higher order" probabilities (e.g., [8]).

If one approves of and has access to such measures, then acceptance can be based on the measure. For instance, successive Π_i 's could be generated by eliminating the next-least probable members of the previous Π . This strategy leads to nested Π_i 's; all decisions would be those mandated by the distribution with the greatest higher-order measure. It would not, in general, be the same as taking an expectation over the expected utility intervals, and ranking the resulting real-values:

$$u(a_\lambda) = \sum_{\forall k \forall i} [\text{Prob}_k(E_i | a_\lambda) u(E_i, a_\lambda)] M(\text{Prob}_k)$$

where M is the higher-order measure. A Bayesian would be committed to the expectation.

Even if one is willing to talk about higher-order probability measures, one should have misgivings about their identification.

There may be uncertainty about the higher-order measure, reflected in some still higher measure. This induces a hierarchy of measures. Presumably the height of the hierarchy is finite. There must be, at some high order, either a determinate measure, or else unmeasured indeterminacy. If the former, then one should be suspicious about the source of a determinate higher-order measure; why is the probability of a distribution certain, but the distribution uncertain? The higher-level is not inherently more robust (note that the order of the sums can be reversed). Just as a small error in a probability can change a decision, so can a small error in a higher-order probability change a decision.

If on the other hand there is unmeasured indeterminacy, the expected-expected utilities will be intervals. This is essentially no different from the interval expected utilities from indeterminate zero-order probabilities.

So acceptance can be conceptually related to higher order probabilities, but is not immediately subsumed or improved by them.

6.3. On Reasonable Conditions.

The proposal requires that all decisions be made with the first Π that is decisive. If all decisions could be made at the highest level of caution, then conditions that do not appear until the less cautious levels are ignored under the proposed decision programme. But what if a condition that any reasonable person would use is ignored? It may seem counter-intuitive that any empirical knowledge whatsoever could be excluded, however reasonable the acceptance of that knowledge, if the agent doesn't need that knowledge to make a decision. But just as there is an upper bound on the caution appropriate for decision analysis, so there is a lower bound. Error in the measurement of utilities and the acceptability of the decision procedure lead to this lower bound [17].

7. Conclusion.

A.I. systems that use interval judgements must sometimes solve partial ignorance decision problems. There are now two approaches. Maximizing expected utility can be followed by lexicographic considerations, or some other secondary criterion. Alternatively, additional assumptions can be made that change probabilities, temporarily, so that maximizing expected utility is sufficient. This paper has discussed the latter approach. Assumptions are accepted in an order that tries to avoid error, and they are accepted only temporarily, for the purposes of decision.

There is still the problem of choosing an acceptance rule, which iteratively generates the next-best assumption. This choice requires considerably more epistemological reflection and practical consideration.

Notes and References.

1. Some have charged that the specification of an interval requires two numbers rather than one; hence, it requires more information. That's silly. Given that some quantity p is in fact 0.67, it follows that p is in the interval [0.34, 0.97]. Intervals are chosen because they offer robust behavior. If practice shows that they are not robust enough, that endpoints matter critically, then future investigators can feel free to use a formalism with indeterminate upper and lower bounds, or with fuzzy sets. Surely one would not revert to point probabilities because they contain "less information."
2. Here, we've taken informativeness w.r.t. decision to be singularity of Π or singularity of the maximal set. Other interpretations of "informative" are possible (such as any restriction of the maximal set to decisions which cannot differ in outcome more than ϵ). These lead to different decision theories.

3. It's also possible to violate nesting when constraints are ordered jointly, and not all constraints are compatible. So if c_1, \dots, c_4 are constraints on Π_1 's, Π_1 may be delimited by $\{c_1\}$, and Π_2 by $\{c_2\}$, and Π_3 by $\{c_4\}$ before Π_3 by $\{c_1, c_2\}$. Π_5 may be delimited by $\{c_1, c_4\}$, which makes sense when $\{c_1, c_2, c_4\}$ is over-determining. If constraints are accepted (rather than knowledge that generates constraint), and acceptance is purely probabilistic, then this kind of situation requires acceptance levels at or below 0.5. With not purely probabilistic acceptance, this situation is more natural.

Note that non-nested Π_i 's would seem irrational via a Dutch Book argument, but the agent still posts consistent odds. For any non-nested convex Bayesian following our proposal, it is possible to construct the behaviorally "equivalent" Π -sequence of a nested convex Bayesian.

- [1] Barnett, J. "Computational Methods for a Mathematical Theory of Evidence," IJCAI 7, 1981.
- [2] Cheeseman, P. "A Method of Computing Generalized Bayesian Probability Values for Expert Systems," IJCAI 8, 1983.
- [3] Dempster, A. "A Generalization of Bayesian Inference," J. Roy. Stat. Soc. 2, 1968.
- [4] Dillard, R. "The Dempster-Shafer Theory Applied to Tactical Fusion in an Inference System," Fifth MIT/ONR Workshop, 1982.
- [5] Fishburn, P. "Analysis of Decisions with Incomplete Knowledge of Probabilities," Op. Res. 13, 1965.
- [6] Garvey, T., Lowrance, J., and Fischler, M. "An Inference Technique for Integrating Knowledge from Disparate Sources," IJCAI 7, 1981.
- [7] Ginsberg, M. "Non-monotonic Reasoning using Dempster's Rule," AAAI, 1984.
- [8] Good, I. *Good Thinking: The Foundations of Probability and Its Applications*, Minnesota, 1983.
- [9] Hurwicz, L. "Some Specification Problems and Applications to Econometric Models," *Econometrica* 19, 1951.
- [10] Kyburg, H. "Conjunctivitis," in Swain, M., ed., *Induction, Acceptance, and Rational Belief*, Dordrecht: Reidel, 1970.
- [11] ———. *The Logical Foundations of Statistical Inference*, Dordrecht: Reidel, 1974.
- [12] ——— and Smokler, H. *Studies in Subjective Probability*, Wiley, 1964.
- [13] ———. "Bayesian and Non-Bayesian Evidential Reasoning," *Computer Science Technical Report 138*, U. Rochester, 1985.
- [14] Levi, I. "Acceptance as a Basis for Induction," in Cohen and Hesse, eds., *Applications of Inductive Logic*, Oxford, 1980.
- [15] ———. *The Enterprise of Knowledge*, MIT, Cambridge, 1980.
- [16] Lopes, L. "Normative Theories of Rationality: Occam's Razor, Procrustes'

- [17] Bed?" (response to Kyburg), *Behav. Brain Sci.* 6, 1984.
- [18] Loui, R. "Decisions with Indeterminate Probabilities," submitted to *Theory and Decision*, 1985.
- [19] Lowrance, J. "Dependency-Graph Models of Evidential Support," Ph.D. Thesis, U. Mass. Amherst, COINS TR, 1982.
- [20] Lu, S. and Stephanou, H. "A Set-theoretic Framework for the Processing of Uncertain Knowledge," AAAI, 1984.
- [21] Nilsson, N. "Probabilistic Logic" SRI Tech. Note 321, 1984.
- [22] Quinlan, J. "Consistency and Plausible Reasoning," IJCAI 8, 1983.
- [23] Shafer, G. *A Mathematical Theory of Evidence*, Princeton, 1976.
- [24] ———. "Constructive Decision Theory," working paper, U. Kansas School of Business Administration, 1985.
- [25] Strat, T. "Continuous Belief Functions for Evidential Reasoning," AAAI, 1984.
- [26] Tong, R., Shapiro, D., Dean, J., and McCune, B. "A Comparison of Uncertainty Calculi in an Expert System for Information Retrieval," IJCAI 8, 1983.
- [27] Wesley, L. and Hanson, A. "The Use of an Evidential Based Model for Representing Knowledge and Reasoning about Images in the Vision System," *IEEE*, 1982.
- [28] ———. "Reasoning about Control: The Investigation of an Evidential Approach," IJCAI 8, 1983.
- [29] ———, Lowrance, J., and Garvey, T. "Reasoning about Control: An Evidential Approach," SRI Tech. Note 324, 1984.
- [30] Zadeh, L. "On the Validity of Dempster's Rule," U.C. Berkeley/ERL M79/24, 1979.