N93-29542-81

161826

p. 10

Decision Analysis
With Approximate Probabilities

Thomas Whalen
Professor of Decision Sciences
Georgia State University
Atlanta, Georgia 30303-3083

ABSTRACT

This paper concerns decisions under uncertainty in which the probabilities of the states of nature are known only approximately. Decision problems involving three states of nature are studied, since some key issues do not arise in two-state problems, while probability spaces with more than three states of nature are essentially impossible to graph.

The primary focus is on two levels of probabilistic information. In one level, the three probabilities are separately rounded to the nearest tenth. This can lead to sets of rounded probabilities which add up to 0.9, 1.0, or 1.1. In the other level, probabilities are rounded to the nearest tenth in such a way that the rounded probabilities are forced to sum to 1.0. For comparison, six additional levels of probabilistic information, previously analyzed in (Whalen, 1991), were also included in the present analysis.

A simulation experiment compared four criteria for decisionmaking using linearly constrained probabilities (Maximin, Midpoint, Standard Laplace, and Extended Laplace) under the eight different levels of information about probability. The Extended Laplace criterion, which was introduced in [Whalen, 1991] using a second order maximum entropy principle, performed best overall.

Risk and Uncertainty

The general problem of decision making under uncertainty involves a set of n states of nature, a set of k alternative actions, and a utility function that assigns a vector of n values to each alterative action; each element of this vector specifies the value of the action under the corresponding state of nature. The k utility vectors typically take the form of row vectors collected into a kXn utility matrix associating a specific value to each (state, action) pair.

Standard treatments of decision making under uncertainty fall into two separate branches: decisions under risk and decisions under ignorance [Resnik 1986]. Under risk, the numeric probability of each state of nature is also assumed to be known or estimated. This enables us to reduce the utility vector of each alternative action to a single number, the expected utility found by adding the product of each utility times the probability of the corresponding state of nature. The action whose expected utility is highest is selected.

Under ignorance, there is no knowledge at all about the probabilities of the states of nature. Various criteria exist for making a decision without recourse to probability. Implicitly or explicitly, each of these criteria replaces the weighting role of the missing

probability values with some other weighting scheme to reduce the vector of possible utilities of an action under the various states of nature to a single value to facilitate comparisons between alternative actions. The Laplace criterion emphasizes all states of nature equally. The Hurwicz criterion (of which maximax and maximin are special cases) emphasizes the most favorable and/or the most unfavorable states of nature. The minimax regret criterion emphasizes the states of nature for which the decision makes the most difference.

Intermediate Cases

In practice, most real decisions use probability information that falls between the well studied extremes of pure risk and pure ignorance. This is especially true in team decision making [Ho & Chu 1972] when one team member assesses a probability distribution but because of time or other constraints can only communicate a standard, concise description of the distribution to the actual decision maker. Each message that can be sent corresponds to a region within a probability space with (n-1) dimensions, where n is the number of states of nature. Note that the authors and publishers of handbooks, almanacs, or other sources of potentially useful information can be viewed as generalized "teammates" of everyone who consults their publications.

For example, sometimes we have enough information to arrange the possible states of nature in order from most probable to least probable, or at least identify some as more probable than others, without being able to numerically specify the probabilities of individual states of nature. This ordinal information may come as a summary message from a teammate, or more directly -- e.g. by observing a random walk process after an unknown number of steps. Alternatively, we may have information about which states of nature, if any, have a probability above a specified threshold.

A very important special case of incomplete probability information arises when probabilities are in rounded form; for example, we may be told that P(A) = .2, P(B)=.3, and P(C) =.4 to the nearest tenth. (A, B, and C are a mutually exclusive exhaustive event set whose unrounded probabilities must sum to 1.) When the probabilities are each rounded to the nearest tenth, it is possible that the sum of the rounded probabilities will not equal 1.0. In practice, rounded distributions of this sort are sometimes communicated as-is, but sometimes the probability distribution as a whole is rounded to the nearest set of three probabilities adding to 1.0. Table 1 shows three sets of exact probabilities, which yield different rounded probabilities when rounded separately but all yield the same rounded distribution when forced to sum to 1.0.

Table 1: Two Methods for Rounding Probabilities

Unrounded Probabilities	Rounded Separately	Rounded to add to 1.0
(.333, .336, .331)	(.3, .3, .3)	(.3, .4, .3)
(.310, .360, .330)	(.3, .4, .3)	(.3, .4, .3)
(.366, .367, .266)	(.4, .4, .3)	(.3, .4, .3)

Linear Probability Constraints & Dempster-Shafer Evidence

The Dempster-Shafer theory of evidence [Shafer, 1976] concerns one particular type of incomplete probability knowledge, represented by basic probability assignments. However, this model does not account for some kinds of probability knowledge that are of great practical importance.

Probability threshold information cannot reliably be expressed by basic probability assignments. For example, with three states of nature we can represent all messages about probability thresholds of 1/4 or 1/3 by basic probability assignments, but not all messages about a probability threshold of 1/2 can be so represented.

When there are only two possible states of nature, the ordinal information that state 1 is more probable than state 2 corresponds to the probability threshold information that $P(s1)>.5$. This can be represented by the basic probability assignment $m(s1)=.5$, $m(s2)=0$, $m(s1Us2)=.5$. However, when there are more than two possible states of nature, ordinal information about probabilities can never be expressed by basic probability assignments.

Rounded probabilities can sometimes be represented by basic probability assignments, but not when the rounded probabilities add up to less than 1.0. For example, probabilities of .33, .33, and .34 would be rounded to .3, .3, and .3. The knowledge that the true probability distribution is somewhere in the region of probability space that rounds to (.3,.3,.3) would provide a useful approximation to the true probabilities, but it cannot be expressed as a basic probability assignment. When probabilities are forced to sum to 1.0, none of the resulting regions of probability space can be represented by basic probability assignments.

All the above cases, and many others, can be expressed by systems of linear constraints on probabilities. In such a case, the available information restricts the probability to lie within a particular region in probability space.

Partial Second Order Ignorance

If a decision maker receives enough information to determine a precise (objective or subjective) probability assessment, the probability region reduces to a single point and the recipient faces a problem of decision making under pure risk. On the other hand, if the recipient can derive no information about the sender's subjective probabilities, the probability region is the whole of probability space, constrained only by the ordinary axioms of probability. In this case, the recipient's problem is equivalent to decision making under pure ignorance.

In the general case, the decision maker knows that the probability distribution over the n states of nature is somewhere within a constrained region r in the probability space. Each point in r specifies an ordinary probability distribution over the states of nature relevant to the original decision problem. This probability distribution together with the payoff matrix for (state-action) pairs in turn specifies an expected value for each action. Thus each point in the region of possible probability distributions specifies an expected utility for each action. The decision maker knows that the true probability distribution over states of nature corresponds to one

of the points in τ, but has no information about the relative likelihood of the points within the region.

This is equivalent to a second order problem of decision making under ignorance. In the second order formulation, the n discrete states of nature are replaced by a continuum of second order "states," where each second order state is a probability distribution over first order states. If the set of second order states includes the full n-nomial probability space, then second order ignorance is equivalent to first order ignorance. In partial second order ignorance, the set of possible second order states equals the region τ (probability distributions that satisfy the constraints arising from partial knowledge about the probabilities).

The payoff for a particular alternative action under a particular second order state equals the expected payoff for that action under the probability distribution over first order states specified by the second order state in question. The decision maker must choose an alternative action in the absence of any information about the second order probability distribution, except that it is within the set of distributions specified by. Thus, it is necessary to rely upon some other consideration to weight the expected return or regret of each probability distribution, in the same way as in ordinary decision making under ignorance.

It is relatively straightforward to find the corner points of a region in probability space defined by a system of linear constraints and to calculate the expected return arising from each alternative action at each corner point. For any possible probability distribution, the expected return for an action is a linear combination of the expected returns of that action at these corner points. Therefore the maximum and minimum expected return for each alternative action can be found by examining only these corner points.

Graphical Analysis When n=3

Suppose that the uncertainty of a decision problem concerns just three possible states of nature. The space of possible probability distributions with respect to these three events forms a planar triangle bisecting the unit cube, as shown in Figure 1. This fact enables us to graph any trinomial probability as a point on a set of triangular coordinates. The three corners of the triangle represent respectively the three trivial probability distributions which assign a probability of 1 to the corresponding states of nature.

Figure 2 shows the 66 regions of probability space that arise from rounding the probability distribution to the nearest decile probability distribution that sums to 1.0. The hexagonal regions represent cases where none of the three rounded probabilities equal zero. The small triangles at the three corners represent the cases when one probability is rounded to 1.0 and the other two are rounded to zero. The pentagons represent cases where one probability is rounded to zero and the other two rounded probabilities are both nonzero.

Figure 3 shows the 166 different regions of probability space that arise from separately rounding each of the three probabilities to the nearest tenth. The hexagonal regions represent cases where the three rounded probabilities add up to 1.0. The small triangles at the three corners represent the cases when one probability is rounded to 1.0 and

the other two are rounded to zero. The trapezoids represent cases where one probability is rounded to zero and the other two rounded probabilities add up to 1.0. The upwards pointing triangles contain probability distributions such as (.86,.06,.08) which when rounded add up to more than 1.0. Finally, the downwards pointing triangles contain probability distributions such as (.84.,03,.13) or (.94,.03,.03) which when rounded add up to less than 1.0.

Decision Criteria

A logical first step in making a decision under uncertainty is dominance screening. Potter & Anderson [1980] discuss dominance screening in the context of linearly constrained Bayesian priors. Ordinary linear programming can find the maximum and minimum values of the difference between the expected utility (EU) of one alternative and that of another. One alternative decision dominates another if the maximum and the minimum difference have the same sign. (A common error is to assume that the maximum EU of the dominated act must be less than the minimum EU of the act that dominates it. In fact two utility ranges can overlap even if one action always has greater EU than the other for each particular feasible probability distribution.)

Typically, more than one nondominated alternative will remain. To reach a final decision, it is helpful to calculate a figure of merit to represent the attractiveness of each action by a single number. When each state's probability is fully determined, expected utility is the figure of merit. When the probability is underdetermined, there are two approaches to calculating a figure of merit. One approach first evaluates the range of expected utilities possible for an action and then reduces this range to a single representative expected utility. The other approach first reduces the range of probability distributions to a single distribution and then calculates just one expected utility using this representative probability distribution.

Representative Utility Approaches

The two most common ways to reduce a range of utilities to a single figure of merit are the maximin criterion and the midpoint criterion. Both are special cases of the Hurwicz family of criteria, which use a general weighted average of the minimum and maximum possible utility: maximin uses a weight of 1.0 for the lower bound and midpoint uses a weight of .5. The maximin criterion expresses conservatism in decision making, while the midpoint criterion seeks to optimize average performance.

The extended Hurwicz criterion selects the action for which
$$\alpha*(\max(E(\text{return}))) + (1-\alpha)*(\min(E(\text{return})))$$
is greatest, where max and min are taken over the set of admissible probability distributions and expectation is taken over states of nature according to each particular distribution. In particular, when the optimism coefficient α equals zero the extended Hurwicz criterion becomes extended maximin. Assuming that the observed decision maker's probability assessment is correct and remains constant for many iterations of the observing decision maker's action, the long-run average return of the extended maximin criterion's selected action cannot possibly fall below the indicated value, while that of other actions might be below this value for some possible probability distribution.

Similarly, when $\alpha=.5$ the extended Hurwicz criterion becomes the extended midpoint criterion, while when $\alpha=1$ it reduces to the extended maximax criterion.

Representative Probability Approaches

On the other hand, many authors [Jaynes, 1968; Gottinger, 1990] argue that uncertainties about probabilities ought to be resolved as objectively as possible; in other words, without reference to utilities. If this principle is accepted, Gottinger has shown that the only reasonable choice for a representative probability distribution from a range is the distribution whose entropy is highest (the Laplace criterion). These arguments are convincing, but their direct application to the probabilities of states of nature can lead to discarding most or all of the available information. For example, the standard maximum entropy (Laplace) form for a complete order over probabilities is equivalent to the maximum entropy form for total ignorance!

This dilemma can be resolved using a second order maximum entropy concept that preserves more real information while satisfying the requirements that motivate the original maximum entropy concept. [Whalen & Brönn, 1990] Rather than considering the probability distribution over the original set of states, we consider a second-order probability distribution over points in probability space (see Figures 1-3). Applying the maximum entropy principle to this distribution implies that all points in probability space should be considered equally likely. Thus the representative point for a region of probability space is the mean point of that region.

Geometrically, the ordinary maximum entropy distribution for a region in probability space (as in Figures 1 & 2) is the point in the region closest to the center of the entire probability space. The second-order maximum entropy distribution for a region is the center of that region itself. Under total ignorance, the region in question is the entire probability space, and both versions of maximum entropy select the same representative point; i.e. the center of the space.

Simulation Experiments

[Whalen, 1991] reports a series of simulation experiments that compared the four methods of determining a figure of merit (Maximin, Midpoint, Standard Laplace, and Extended Laplace) using six different information systems:
(1) the null information system in which the decision maker has no information about probability,
(2) an ordinal information system in which the decision maker can rank the 3 probabilities from lowest to highest (6 possible messages),
(3) an information system that informs the decision maker which probability, if any, is above .5 (four possible messages),
(4) an information system that informs the decision maker which probability, if any, is above 1/3 (6 possible messages),
(5) an information system that informs the decision maker which probability, if any, above .25 (7 possible messages), and
(6) the perfect information system in which the decision maker knows the exact probabilities of the three states.

Ten thousand trinomial distributions were generated according to a uniform second-order distribution: $p1 = 1-R^{\frac{1}{2}}$, $p2 = S*(1-p1)$, $p3 = 1-p1-p2$ where R and S are uniformly distributed random fractions. Ten thousand 3X3 utility matrices were randomly generated; the highest utility in each matrix was 100 and the lowest zero, with other utilities uniformly distributed. Each pairing of a criterion with an information system selected an action, and the expected utility of that action was recorded for a total of ten thousand iterations. The lowest mean expected value was 64.255 (maximin criterion, null information system), and the highest mean expected value was 71.748 (perfect information system).

In the present research, the same benchmark set of 10,000 probability distributions and utility matrices was used to examine the performance of the decision criteria using the richer information provided by probabilities rounded to the nearest tenth. The label "Round:1.0" refers to the information system in which rounded probabilities are forced to sum to 1.0, while the "Round:.9-1.1" label refers to the information system which rounds the three probabilities separately. For these two information systems, a fifth decision criterion is also shown; in this criterion, the expected value is simply calculated using the three rounded probabilities. (In the "Round:.9-1.1" system, rounded probabilities are used without regard to whether they sum to 0.9, 1.0, or 1.1.)

Table 2 summarizes the findings of [Whalen, 1991] together with the new experiment (the rows labeled "Round:.9-1.1" and "Round:1.0"). The table shows the mean expected utility of each combination of one of the seven information systems with one of the four decision criterion, expressed as a percentage of the range of mean expected utility from the lowest to the highest; 0% means the lowest observed utility (64.255) and 100% means the highest observed utility (71.745). Thus, the percentages represent the proportion of the maximum benefit that can be derived from probability information.

TABLE 2

	# of Messages	Standard Laplace	Maximin	Midpoint	Extended Laplace	As Rounded
None	(1)	48.0%	0.0%	33.9%	48.0	
Ordinal	(6)	48.0%	81.1%	89.7%	88.6%	
Threshold=1/2	(4)	80.9%	78.0%	86.4%	88.6%	
Threshold=1/3	(6)	48.0%	84.7%	92.4%	92.2%	
Threshold=1/4	(7)	79.0%	85.2%	91.6%	92.3%	
Round:1.0	(66)	95.8%	97.7%	98.56%	98.57%	98.47%
Round:.9-1.1	(166)	98.6%	98.8%	99.1%	99.5%	99.4%
Perfect	(10000)	100.0%	100.0%	100.0%	100.0%	

Several interesting observations can be made based on these results. Not surprisingly, there is a general tendency for the performance of the various techniques to increase with increasing richness of information as measured by the number of alternative messages. But there are some noteworthy exceptions.

The Ordinal information system always leads to poorer performance than the probability threshold 1/3 even though both have six messages; furthermore, in the two representative probability approaches (Standard Laplace and Extended Laplace), the six-message Ordinal information system is actually inferior to the four-message information system with probability threshold .5! Under the Midpoint criterion, the seven-message information system with threshold .25 is inferior to the six-message information system with threshold 1/3, while under the Standard Laplace criterion the four-message information system with probability threshold .25 outperforms both six-message information systems and the seven-message information system. The only decision criterion which comes close to consistently rewarding richer information with better performance is the Extended Laplace, although even here the performance with ordinal information is very slightly poorer than the performance with information based on a probability threshold of .5.

Comparing decision criteria under a given information system, the Extended Laplace consistently outperforms the others except in the case of the Ordinal information system, in which it is not quite as good as the Midpoint criterion. Despite strong theoretical endorsements (Jaynes, 1968; Gottinger, 1990), the Standard Laplace is consistently the worst except in the case of the information system with probability threshold = .5, in which it is better than the maximin criterion. These results seem to imply that the Extended Laplace is the correct way to apply the principle of maximum entropy to problems of this type.

The relationships among the decision criteria are summarized in Figure 4 for the three probability threshold information systems and the two rounded probability information systems. (The horizontal axis, labeled "bandwidth," is the logarithm to the base 2 of the number of messages in the information system, ranging from 2 bits for the four-message system to 7.375 bits for the 166-message system.)

REFERENCES

Gottinger, H. W. (1990), Decision problems under uncertainty based on entropy functionals, Theory & Decision 28:2 p.143-172.

Ho, Y-C. and Chu, K.C. (1972) Team decision theory and information structure in optimal control problems (Parts I & II). IEEE Trans Autom. Control AC-17, 15-38.

Jaynes, E. T. (1968), Prior probabilities, IEEE Trans Systems Science & Cybernetics SSC4:3 p.227-241.

Potter, J. & Anderson, B. (1980), Partial prior information and decisionmaking, IEEE Trans. on Systems, Man, and Cybernetics SMC10:3 p.125-133.

Resnik, M. (1987) Choices: An Introduction to Decision Theory. U. of Minnesota Press, Minneapolis, Mn.

Shafer, G. (1976), A Mathematical Theory of Evidence, Princeton Univ. Press.

Whalen, T. & Brnn, C. (1990), Hurwicz and regret criteria extended to decisions with ordinal probabilities, Proc. 1990 NAFIPS Conf. p.219-222.

Whalen, T. (1991), Decisions With Underdetermined Probabilities, Proc. 1991 NAFIPS Workshop p.8-12.

Figure 1
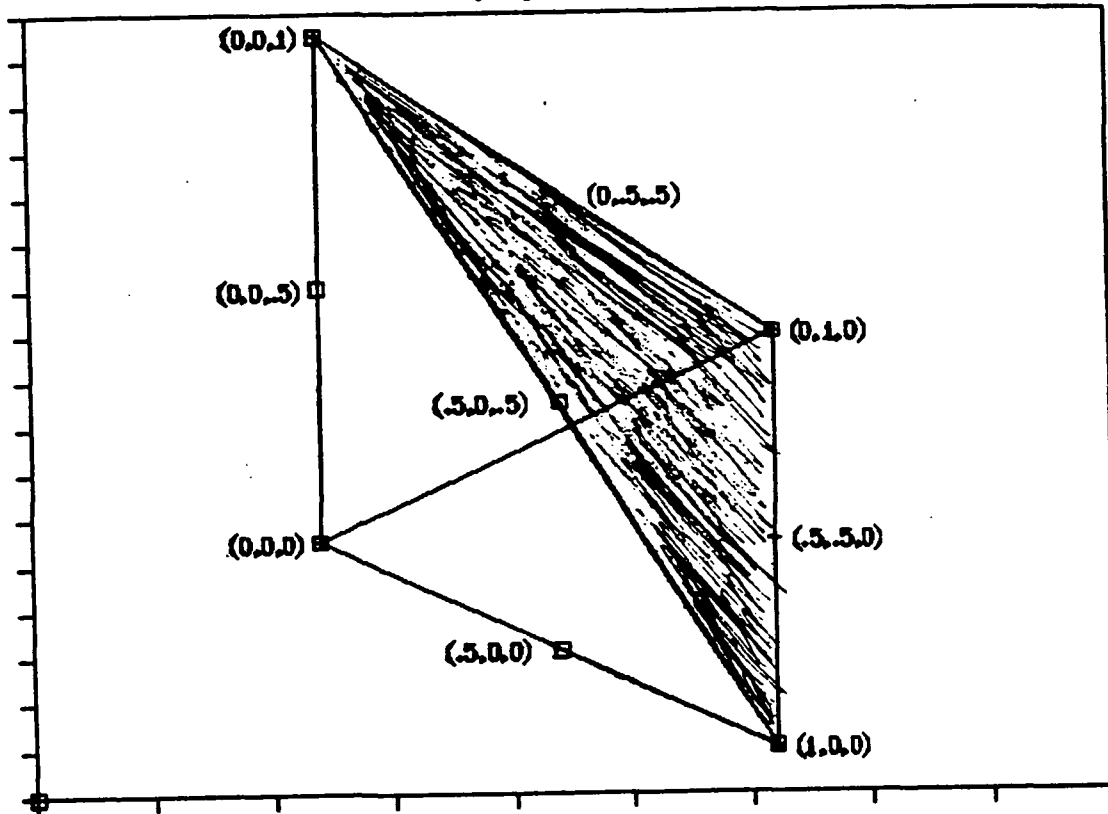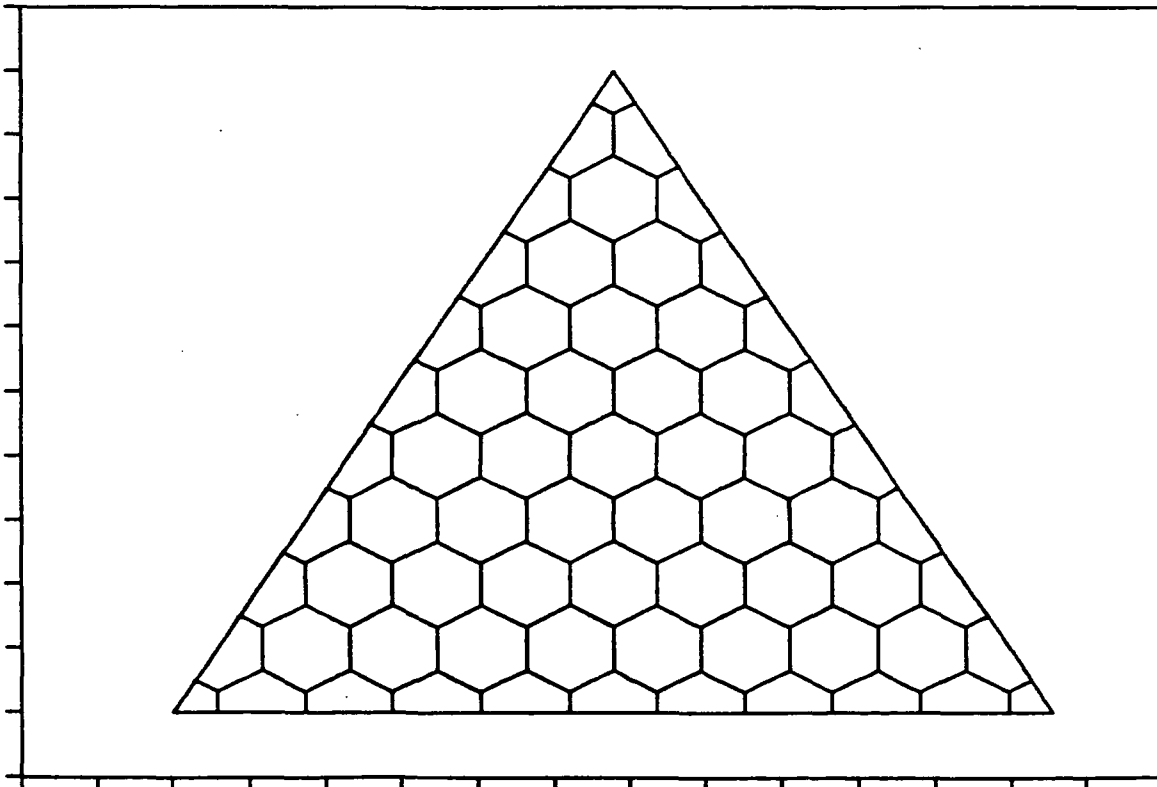### Probability Space (n = 3)

Figure 2
Decile Probabilities Summing to 1.0

Figure 3
Probabilities Rounded to Nearest Tenth

Figure 4
Results of Simulation Experiments