AD-770 583

SOME CRITICISMS OF THE GENERAL MODELS USED IN DECISION MAKING EXPERIMENTS

Dirk Wendt

Michigan University

011313-8-T

THE UNIVERSITY OF MICHIGAN

AD 770583

ENGINEERING PSYCHOLOGY LABORATORY

Technical Report

Some Criticisms of the General Models Used in Decision Making Experiments

DIRK WENDT

Sponsored by:

Advanced Research Projects Agency
Department of Defense

DDC
DEC 6 1973
RECEIVED
B

Monitored by:

Engineering Psychology Programs
Office of Naval Research
Contract No. N00014-67-A-0181-0049
NR 197-021

Approved for Public Release; Distribution Unlimited

Reproduction in whole or in part is permitted for any use of the U.S. Government

Administered through: 2 November 1973

OFFICE OF RESEARCH ADMINISTRATION · ANN ARBOR

40

SOME CRITICISMS OF THE GENERAL MODELS

USED IN DECISION MAKING EXPERIMENTS

Technical Report

2 November 1973

Dirk Wendt

Engineering Psychology Laboratory

The University of Michigan

Ann Arbor, Michigan

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER 011313-8-T	2. GOVT ACCESSION NO	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) SOME CRITICISMS OF THE GENERAL MODELS USED IN DECISION MAKING EXPERIMENTS		5. TYPE OF REPORT & PERIOD COVERED Technical
		6. PERFORMING ORG. REPORT NUMBER None
7. AUTHOR(s) Dirk Wendt		8. CONTRACT OR GRANT NUMBER(s) N00014-67-A-0181-0049
9. PERFORMING ORGANIZATION NAME AND ADDRESS Engineering Psychology Laboratory Institute of Science & Technology University of Michigan Ann Arbor, Michigan 48105		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS NR 197-021 ARPA Order No. 2105
11. CONTROLLING OFFICE NAME AND ADDRESS Advanced Research Projects Agency 1400 Wilson Boulevard Arlington, Virginia 22209		12. REPORT DATE 2 November 1973
		13. NUMBER OF PAGES 39
14. MONITORING AGENCY NAME AND ADDRESS (if different from Controlling Office) Engineering Psychology Programs Office of Naval Research, Department of the Navy Arlington, Virginia 22217		15. SECURITY CLASS. (of this report) Unclassified
		15a DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Approved for public release; distribution unlimited.

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Risky decisions
Maximization model
Normative model
Probability cues

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

The general normative model of expectation maximization is outlined, and criticized for several reasons. It may not be appropriate as a normative model in a variety of situations where it is assumed to be rational. Some of its conditions, e. g., independence of evaluation-of-aspects and probability-revision cues, and correctness of the simple additive utility model, may not be met. Moreover, deterministic models may be too strong to predict human behavior properly. Perhaps they should be replaced by probabilistic ones. The emphasis of this paper, however

20.

is not to doubt the applicability of the model in principle but rather to point at some problems where more research is needed.

Some Criticisms of the General Models
Used in Decision Making Experiments

Introduction

In the beginning, there was the idea of expectation maximation as a
normative model in decision making: given a choice of several courses of
action $\underline{a}_i$ you should always select that alternative for which the expected
value

$$\underline{EV}_i = \sum_j p(s_j) \, v \, (a_i, \, s_j) \tag{1}$$

is largest, where $p(s_j)$ is the probability of some other event s_j which you
cannot control, and $v(a_i, \, s_j)$ is the (monetary) outcome for you if you choose
a_i, and the uncontrollable event s_j comes up. Decision situations to which this
model is applicable are conveniently represented by payoff matrices.

$(v(a_i, \, s_j))$: state of nature, or other event not under
DM's control

choice of action
(controlled by DM)

	s_1	s_2	$\cdots$	s_j	$\cdots$	s_m
a_1				$\vdots$		
a_2				$\vdots$		
$\vdots$				$\vdots$		
a_i	$\cdots\cdots\cdots\cdots\cdots\cdots$			$v(a_i, \, s_j)$		
$\vdots$						
a_n						

Bernoulli (1738) was the first (in literature) who found out that this model does not work, either in betting behavior or in insurance buying. He tried to get around it by replacing monetary values with utilities, assuming a negatively accelerated utility function. Many other authors have followed him in recent years (about since the 1950s) assuming all kinds of utility functions, even positively accelerated ones for gamblers, and multi-attribute utility functions for more complex situations. (For an excellent review of the literature on the latter approach, see e.g. Fischer (1972)).

Von Neumann and Morgenstern (1944) have re-emphasized expectation maximization as a normative model. Experiments following their book have shown that not only values but also probabilities need some subjective distortion to make the model descriptive. As a combination of these transformations, Edwards (1954) proposed the subjective expected utility maximization model, SEU, when both values $v(a_i, s_j)$ and probabilities $p(s_j)$ are replaced by their subjective transforms, but the principle of expectation maximization is preserved.

Coombs and collaborators (from Coombs and Pruitt (1960) up to the most recent, not-yet-written-up papers) enriched the model by an additional risk component. They assumed that risk taking, for most Ss, has a subjective value in itself, leading to an ideal (i.e. most preferred) level of risk for any S, implying a single peaked preference function over risk. In most cases, risk was somehow related to the variance of the possible outcomes of a choice of action. Coombs called his risk-expectation approach Portfolio Theory.

Both Edwards' SEU and Coombs' Portfolio approach were mainly intended to

make the original normative expectation maximization model more descriptive, more predictive of actual behavior.

However, the expectation maximization approach has competitors, also as a normative model (for a review of these alternative normative decision criteria see, e.g., Schneeweiss (1966), or Wendt (1970)). But, so far, only little use has been made of these alternative criteria in the experimental literature. The exception is, the minimax criterion, which has been frequently used in competitive experimental zero-sum games where s_j, the event not under DM's control, is the decision of another intelligent $\underline{S}$ trying to minimize DM's payoff to maximize his own.   In this case, it is advisable to decide in such a way that the worst possible outcome is still the smallest loss for DM, in the given situation. This calls for a strategy where DM first finds the minimum payoff in each row of the payoff matrix (because that is the column his opponent will like most to select), and then picks the row a_i in which this minimum over columns is largest. Because of this procedure--search for minima within rows first, and for row with maximal minimum then--this strategy has been called the Minimax Strategy. It was first defined by Wald (1950).

Savage (see 1954, though the idea was proposed earlier) refined this strategy a rather interesting way. If the events s_j not under his control are very different in attractiveness to DM, his main concern should be how much he loses by making the wrong decision in a given situation. Thus, Savage considered it appropriate to deduct from each entry in the payoff matrix the most attractive value

(maximum) in its column, leaving in each cell the amount DM would get less than this maximum if the respective column s_j obtained, and he had not chosen the row with the maximum payoff in that column. These quantities, $r(a_i, s_j) = v(a_i, s_j) - \max_i [v(a_i, s_j)]$, have been called regrets, and the minimax regret strategy advises DM to apply the minimax principle (as stated above) to regrets $(r(a_i, s_j))$ rather than actual payoffs $(v(a_i, s_j))$.

The following example is designed to show that expectation maximization minimax payoff and minimax regret strategies can lead to different choices of action:

	payoffs $v(a_i, s_j)$		regrets $r(s_i, s_j)$			recommended choice by criterion
	s_1	s_2	s_1	s_2	EV	
a_1	-10	+20	-110	0	+14	expectation maximization
a_2	+2	-1	-98	-21	-.4	minimax payoff
a_3	+100	-50	0	-70	-20	minimax regret
$\max_i [v(a_i, s_j)]$	+100	+20	← ↑ ← ↑ enter the calculation of ↑ ←			(These all are considered normative rational strategies.)
$p(s_j)$	.2	.8				

[3]

In this example, the argument for minimax regret goes as follows: according to the minimax payoff criterion, DM should choose a_2 which secures him at least -1 (in the worst case s_2, given he has chosen a_2) whereas with choice a_1 or a_3 he could lose 10 or 50, respectively. However, if s_1 comes up and he

has chosen a_2 as recommended by the minimax payoff strategy, he gets 98 less than what he could have made if he had chosen the best action under assumption that s_1 obtains, which would be a_3. If he chooses a_3, and s_2 comes up, he has to regret only 70 which he would have got more if he had chosen the optimal action a_1 under assumption of s_2. (Under the expectation maximization criterion, expected payoff maximization and expected regret minimization always lead to the same choice of action.)

If applied in risky decision situations where the event s_j is not under another intelligent being's control but rather random, the minimax criteria are rather pessimistic or cautious strategies since they always assume that the worst possible event s_j can obtain whatever DM's decision a_i might be.

This is considered adequate in games against a hostile opponent but not when the s_j are random events in nature. Then, expectation maximization is recommended in most cases.

Therefore, the normative model of expectation maximization will be outlined in a little more detail before going into the criticism that is the topic of this paper. As indicated, the model prescribes the choice of that action a_i for which the quantity

$$\underline{EV}(a_i) = \sum_j p(s_j) \; v(a_i, \; s_j)$$

[4]

is maximal over all a_i. $p(s_j)$ is considered DM's personal or subjective probability that s_j will obtain. $p(s_j)$ obeys the axioms of probability theory, and can be modified according to the rules of Bayes' theorem if relevant data d_k bearing on s_j are observed:

$$p(s_j|d_k) = p(s_j) \cdot p(d_k|s_j) \; / \; \sum_j p(s_j) \; p(d_k|s_j)$$

[5]

where $p(s_j|d_k)$ is the posterior probability of s_j given the datum d_k,

$p(s_j)$ is the prior probability of s_j, and $p(d_k|s_j)$ is the likelihood, or probability of d_k given s_j.

If there are several independent data d_k relevant to the distribution over s_j, their respective probabilities are multiplied together to reveal the overall impact of the data on prior probabilities:

$$p(s_j|d_1, d_2, \ldots, d_r) = \frac{p(s_j) \underset{k}{\pi} p(d_k|s_j)}{\underset{j}{\Sigma} p(s_j) \underset{k}{\pi} p(d_k|s_j)}$$

[6]

Thus, with the use of data, we have

$$\underline{EV}\ (a_i) = \underset{j}{\Sigma} \frac{p(s_j) \underset{k}{\pi} p(d_k|s_j) \ v(a_i, s_j)}{\underset{j}{\Sigma} p(s_j) \underset{k}{\pi} p(d_k|s_j)}$$

[7]

In most more complex real life interactions, no monetary payoffs $v(a_i, s_j)$ are known. Instead, many different aspects or dimensions may enter the evaluation of the various outcomes of decisions.

Several solutions have been suggested for this evaluation problem; for a recent review of these see Fischer (1972). Most of them assume an additive combination of the various aspects to be most adequate, proposing a weighted sum or average of the attributes x_h

$$v(a_i, s_j) = \underset{h}{\Sigma} w_h x_h (a_i, s_j)$$

[8]

where w_h is the weight, or importance of the h-th dimension or attribute of the situations, and $x_h (a_i, s_j)$ is the score of the situation (a_i, s_j)

on the h-th dimension, i.e., the amount to which the situation (a_i, s_j) possesses the attribute h. As Keeney 1972 has shown, this additive model is the limiting case of the more general multiplicative model,

$$1 + W\, v(a_i, s_j) = \pi_h \left[1 + W\, w_h\, x_h(a_i, s_j)\right] \tag{8a}$$

such that

$$v(a_i, s_j) = \frac{\pi_h \left[1 + W\, w_h\, x_h(a_i, s_j)\right] - 1}{W} \tag{8b}$$

With insertion of [8] into [7], we get

$$\underline{EV}\,(a_i) = \sum_j \frac{p(s_j)\, \pi_k\, p(d_k|s_j)\, \sum_h w_h\, x_h\,(a_i, s_j)}{\sum_j p(s_j)\, \pi_k\, p(d_k|s_j)} \tag{9}$$

In this equation, variables and parameters can be categorized into three classes:

some are rather stable parameters or characteristics of DM, like w_k;

some are rather stable characteristics of DM's environment, like

a_i, s_j, $p(s_j)$, $p(d_k|s_j)$;

and some have to be taken into account in the evaluation of each

particular situation, like $x(a_i, s_j)$, and observation of d_k.

This third category contains those cues or characteristics of the situation DM looks at to make up his mind: he observes the status of data d_k relevant to his revision of probabilities $p(s_j)$ to $p(s_j|d_k)$, and ne estimates the attribute scores $x_h(a_i, s_j)$ to get utility evaluations $v(a_i, s_j)$. The normative model above considers these two kinds of cues, d_k and $x_h(a_i, s_j)$ to be very distinct classes of variables which play completely different roles in reaching a decision (choice of a_i).

Assuming that Eq. [9] is the optimal rational strategy to select an action a_i, it is easily seen that human beings might have a hard time doing all these calculations in their heads, and then come up with the right decision. It is rather obvious that they need some help.

This paper is intended to make some criticism of the approach used so far, i.e. to question if the technologies developed and tested in experimental decision making research really can help the decision maker in real life situations. To be helpful in this sense at all, it would be necessary

(a) that the decision maker actually agrees that the models proposed (as outlined above) can be considered as rational normative models which meet his needs and goals

(b) that these models are descriptive enough of human decision behavior so that it is possible for the decision maker to do what the model prescribes.

More particularly, we will discuss four major points:

first, we question the assumption of expectation maximization as a rational concept at all;

second, we cast some doubt on the separation of probability-estimation and outcome-evaluation cues (d_k and x_h, respectively) in the judgment of real life situations;

      <u>third</u>, we question the weighted-average multi-attribute utility
         model in particular; and

      <u>fourth</u>, we suggest the use of probabilistic models in decision
         making (just as we do in other fields of psychology) rather
         than deterministic models (as it has been done in empirical
         studies of risky decision making so far.)

I am aware of the fact that all these criticisms have been made before -
since "there is no new thing under the sun" anyway (The Bible, Ecclesiastes
1:9). But nobody has listened, so far. What I leave out is questioning the
generalizability of experimental laboratory results to real life situations--
this problem, with respect to decision making, has been reviewed recently by
Winkler & Murphy, 1973.

All this has been written up not to point out that our current decision
technology is useless--on the contrary, it has made tremendous progress
during the last twenty years--but rather to show that it still is far from
perfect, and that much more research is needed.

Criticism of expectation maximization as a normative model

Maximizing expectation in risky decision situations may become quite fatal for DM if high losses are possible which can throw DM out of the business. Actually, the expectation maximization strategy should not be recommended to any DM unless he is infinitely rich, so that he can afford occasional losses no matter how high they might be.

Bernoulli (1738) tried to get around this problem by assuming non-linear, negatively accelerated utility functions dependent on DM's total wealth as a normative model for insurance buying. This did not solve the problem satisfactorily, either, but in principle, nonlinear utility functions can do it by assessing an infinitely large negative utility to bankruptcy. However, such a utility function would violate the Archemedian axiom which is necessary for any utility function. [DeGroot, 1970, p. 102]. The classical gambler's ruin problem has been discussed in Ch. XIV of Feller (1968), showing that even with unfair bets (i.e., with odds less than one for the DM) you can expect to gamble fairly long without being ruined; you even have a considerable probability of gain (See Table 1 from Feller, p. 347).

TABLE 1

ILLUSTRATING THE CLASSICAL RUIN PROBLEM

p	q	z	a	Probability of Ruin	Probability of Success	Expected Gain	Expected Duration
0.5	0.5	9	10	0.1	0.9	0	9
0.5	0.5	90	100	0.1	0.9	0	900
0.5	0.5	900	1,000	0.1	0.9	0	90,000
0.5	0.5	950	1,000	0.05	0.95	0	47,500
0.5	0.5	8,000	10,000	0.2	0.8	0	16,000,000
0.45	0.55	9	10	0.210	0.790	−1.1	11
0.45	0.55	90	100	0.866	0.134	−76.6	765.6
0.45	0.55	99	100	0.182	0.818	−17.2	171.8
0.4	0.6	90	100	0.983	0.017	−88.3	441.3
0.4	0.6	99	100	0.333	0.667	−32.3	161.7

The initial capital is z. The game terminates with ruin (loss z) or capital a (gain $a - z$).

Normative models for gambling under limitations of time have been discussed by Dubins & Savage (1965). They do not maximize expected value, but rather maximize the probability of ending up with a prespecified total amount. Expectation maximization might serve only as an approximation to these strategies (and it would be well worth while studying how good these approximations are).

These facts have probably led to the widespread acceptance of the expectation maximization model it has gained now.

For any DM with limited wealth, the minimax strategy (or at least a part of it) might be more advisable in some situations.

Although a pure minimax strategy seems to make sense from a risk-avoidance point of view, it would not give DM a chance to make much of a fortune. It is defensive rather than designed to make DM successful.

What is called for is a decision strategy which is partly defensive like the minimax criterion to avoid bankruptcy, and partly prosperous like the expectation maximization criterion to enable DM to have some success. Such a criterion has been proposed by Hodges and Lehmann (cf. Schneeweiss 1966) who suggested to weight both expectation $\sum_j p(s_j) \, v(a_i, s_j)$ and the worst outcome $\min_j [v(a_i, s_j)]$ by a confidence parameter β and $(1-\beta)$, respectively: DM is required to maximize the quantity

$$\underline{Z}(a_i) = \beta \, \underline{EV}(a_i) + (1-\beta) \, \min_j [v(a_i, s_j)] \tag{10}$$

with $\underline{EV}(a_i)$ as specified in Eq. [1] and [9].

Formally, this means just giving some additional weight $(1-\beta)$ to the worst outcome (which is contained in $\underline{EV}(a_i)$ weighted by its probability anyway). This makes the model more cautious than plain $\underline{EV}$ maximization but still may do a rather poor job as long as β is considered a rather rigid parameter. What might help, is to make it more adaptable, i.e. to adjust the amount of pessimism $(1 - \beta)$ to DM's actual situation. For instance, β should take into account DM's current wealth or total capital, to adapt to the extent to which DM can afford expectation maximization.

Thus, instead of a rigid constant parameter, β should be a (antitone) function of DM's current total capital y, and a monotone function of $\min_j[v(a_i, s_j)]$ itself, and maybe respect some other aspect of the distribution over outcomes, too.

However, giving an extra weight to the worst outcome may not be the only possible improvement of the expectation maximization model. Other aspects of the distribution $(p(s_j|d_k))$ over possible outcomes (a_i, s_j) may be important, too, as Coombs assumes in his Portfolio Theory, and as Pollatsek and Tversky 1970 assume in their theory of risk which considers risk as a linear combination of expectation and variance. But they are more concerned with empirical findings in the perception of risk, and less with normative models and decision aids.

Another point open to discussion is whether all this should be applied to regrets rather than to utilities (payoffs). This distinction has been neglected in the so far predominant expectation maximization model where it is of no practical importance because maximization of expected value leads to the same choice of action as does maximization of expected regret. This does not hold if we consider a minimax strategy, or a mixture of minimax and expectation maximization strategies, as shown in the example on p. 4.

The transformation of payoffs into regrets is quite a meaningful one considering its implications for DM's choice of action. It concerns the exclusion of irrelevant portions of the payoffs, i.e., consequences of the outcomes which the DM cannot influence by his choice of action in any way.

In a sense, this corresponds to Restle's 1961 modification of the BTL choice model. DM should not make his choice dependent on aspects of the outcomes which he gets in any case, i.e. with any choice which is the intersection of the alternatives in Restle's riskless choice model, or here: DM should not make his choice dependent on aspects (utilities) of outcomes which he could not have gotten with any choice because the state of the world allowing these outcomes did not obtain. Please note that I do not pretend that Restle's deduction of the intersecti n from the numerator in the choice probabilities corresponds to the deduction of column maxima in the calculation of regrets - I am just pointing out that the motivation behind both procedures is similar.

Since the expectation maximization model has been so predominant in experimental psychological research for years, very little has been done exploring the explanatory value of the alternative normative models discussed above.

The over-exaggeration of cupidity motivation

The normative models generally applied in decision making experiments hinge upon the assumption that man is motivated by cupidity. Thus, experimenters build their whole system around this assumed motivation, and worry about the flatness around the maxima of the payoff functions because they feel that subjects won't try hard enough since they lose little by being non-optimal (v. Winterfeldt & Edwards, 1972).

Cupidity may be indeed a strong motivation for human subjects, and some recent experiments in decision making with considerable amounts of money at stake have shown that payoffs actually do make a difference in subjects' behavior (Snapper, 1973; Areen, 1973).

On the other hand, Fryback, Goodman, & Edwards (1972), show that payoff magnitude may not be really as important as other factors. Thus, monetary payoffs are not that important under all circumstances. On the contrary, in more real-life like situations, just the type of subjects we mostly use in our laboratory decision making experiments, i.e. students, show that they are much more motivated by ambition than by cupidity. E.g., in athletic competitions, the differences in scores are usually very small between the top competitors, so that we really can talk about a "flat payoff function"--but the smaller the differences get, the harder they try to get to the top. For them, rank order position counts much more than plain score, and in most cases

there is no money involved at all (except in those cases where they compete
for athletic scholarships). Most students have this ambitious attitude not
only in athletic but also in intellectual tasks, and thus try hard to be
optimal even in decision making experiments with little or no money at stake--
provided they get another kind of feedback, like comparison with their peers.
A normative model which accounts for this kind of motivation should be based on
some measure like the probability of becoming first, or more generally, of
moving up in the rank order.

Thus, a rank order transformation plus some ambition (rather than cupidity)
motivation may become a remedy for the flat maximum problem. I have the
impression that some companies have been using such techniques for some time
to stimulate their employees without spending considerably more on salaries,
although we should wonder about the moral implication of such practice.

The cost of precision

Another point should be made discussing the flat maximum problem: the
normative expectation maximization model does not take into account any cost of
precision, i.e. how much effort it might cost a human subject to find the
absolute optimum precisely. To explain this by an example involving no
probabilities (just to make it simpler) assume a customer wants to buy some
standard item he can mail order from one of several catalog stores. Of
course, he wants to buy it at the lowest price possible (since the quality
is the same in all stores), and going through the catalog of some stores he
finds that the prices for that particular item vary some cents. Now, how many
catalogs should he go through before he decides to order it? Since he does
it in his spare time, where time spent searching does not "cost" him anything,
moneywise, he should (normatively) check out all catalogs he can get, and make
a couple of phone calls to local stores in addition. However, he will soon
feel that the small advantage in price he might find is not worth all that fuss,
and thus he probably makes his order after looking at only two or three of the
catalogs available to him. The subject in a laboratory decision making experiment
may be in a similar situation: he may know that he could make some pennies more
if he thought it over more carefully, but he also wants to get home, eventually.
I wonder if anybody has ever designed an experiment to find something about this
time/payoff trade-off. Of course, in such an experiment we would have to intro-
duce some utility function for time (which could be analogous to sampling costs
in previous experiments). Such an experiment would be the natural consequence
of the v. Winterfeldt & Edwards flat maximum paper.

The non-persistence of monetary value

The normative model of decision making, and most experiments applying it,
assume a persistency of the monetary values involved as payoffs. However, in
real life (especially in the professional lives of business executives for whom
many of the normative models are made) money has no persistent value at all;
it rather varies considerably with, e.g, the time and terms of its availability.
Any normative model that is really meant to be helpful for actual decision making
should incorporate these dynamics. The extreme case to observe the effects of
such non-persistence of monetary value on decision making would be the inter-
national money traders who make profits out of small differences in exchange
rates for foreign currencies. But also non-professionals in this field do have
preferences about when to pay or receive what - the whole credit card business
lives on that.

Criticism of the assumption of independence of
probability and utility cues

The normative model of expectation maximization assumes (in Eq. [9]) independence of cues $x_h(a_i, s_j)$ contributing to the evaluation of possible outcomes $(a_i, _j)$, and cues d_k leading to the revision of probability distributions $p(s_j|d_k)$ over possible states (s_j).

I doubt if this separation of parameters really makes sense, either normatively or descriptively.

In the more complex situation or real life where you have different sources of information delivering cues on the state of the situation, it may depend just on the arbitrary definition of objectives whether a cue is an aspect of utility, or a datum for probability revision. An example:

Does the fact of having an air conditioner in a car increase its utility as a utility aspect x, or does it rather increase the probability $p(s)$ of the particular state s that you will feel more comfortable in it when it is hot? Let us elaborate this example in a little more detail. In the first stage, let us assume you have the choice of buying either a car with an air conditioner, or a car without it. The former will be little more expensive than the latter - so it depends on the usefulness or utility of the air conditioner, if it pays for you to pay these extra expenses. This, again, will depend on how much you are bothered by heat, and this, in turn, will depend on how much heat there is, i.e. on the weather condition, and on the fact of

having an air conditioner or not. Thus, your final utility for the whole
situation (at the end of the decision tree) will depend (among others) upon
the cues: "weather situation", "air condition available or not", "DM's
susceptibility to heat", "acceptable retail price when selling it", "price
paid for car", etc. At the same time, these very same cues influence the
conditional probabilities of the various outcomes (branches) of the decision
tree which are partly contingent on preceding conditions: "feeling confortable"
is dependent on weather conditions and availability of air conditioning. Also,
having air conditioning increases the probability of a good resale price -
because it will be of potential value to a prospective buyer. Other cues are
independent, like weather conditions and the kind of car you buy. In some cues,
direct dependencies between utility and probability cues are obvious, like the
relation between "probability of getting an acceptable retail price" and the
amount you consider an "acceptable retail price" which might be a kind of a
psychometric function.

Thus, many cues have bearing on both utility and probability aspects, and
in many cases is just a question of truncation of the decision tree if they
enter either of them or both.

Practically every decision tree has to be truncated somewhere - we cannot
follow all of its ramifications into eternity, e.g. we would not care to consider
how our car buying decision today may affect the life of our great-grand children,
although it is quite clear that it can influence them a lot, probability-wise.
In the example above, we might have truncated the decision tree before con-
sidering the possibilities of selling the car after a while. Then, we would

not have been concerned with "probability of getting an acceptable retail price", but this probability (and the price we consider acceptable) would certainly be contained implicitly in the "value of the car" we are left with. On the other hand, by extending the decision tree, we could consider almost all cues as data for modifications of conditional probabilities.

In cases like this, where probabilities and utilities of particular attributes can, ideally, substitute for each other, we might wonder what the tradeoff function between these measures are. Of course, in all individual cases, this will depend on the attribute measured, but in general we might think of such functions as psychometric functions in psychophysics where we have some physical quantity on the abscissa, and a probability (of noticing the physical quantity) on the ordinate. Actually, such a technique could be applied in the used car example above if we plot the retail price on the abscissa, and the decision maker's subjective probability of obtaining this retail price on the ordinate.

An example of using probabilities of attaining some desirable outcome (i.e. victory in a battle) rather than utilities or monetary values in a payoff matrix is given in Coombs, Dawes & Tversky 1970, p. 208.

In real world examples, it may be hard to find cues which cannot be interpreted in these two different ways. In many cases, however, it may be possible to restructure the situation such that the ambiguity disappears, dependent on how you re-define the possible states s_j (which are arbitrary in most cases, anyway), and where you truncate the decision tree. Such restructuring may solve the problem of independence between cues for probability revision and

for utility assessment formally in many cases, but leaves us with the problem of how the subjects perceives these cues, and uses them in his decisions. (Luce & Krantz (1971) in their axiomatization of conditional expected utility, however, circumvent this problem by going up the decision tree rather than down, i.e. by considering utilities for whole gambles (i.e. larger branches of the decision tree) rather than single outcomes.)

These necessary re-definitions of the set (s_j) of possible states of nature may imply that all games be in normal form, that more complex decision trees be reduced to complete preprogrammed but adjustable strategies of the opponent (or nature).

The interchangeability of cues for probability revision and utility assessment throws some light upon the importance of independence between evaluation scores $x_h(a_i, s_j)$, and data probabilities $p(d_k|s_j)$: in Eq. [9], we multiply the factors $p(s_j)\pi_k p(d_k|s_j)$ (which are constants with respect to the subscript h) into the multi-attribute utility function, ending up with

$$\underline{EV}(a_i) = \sum_j \frac{\sum_h w_h\, p(s_j)\pi_k p(d_k|s_j)\, x_h\,(a_i,\, s_j)}{\sum_j p(s_j)\pi_k p(d_k|s_j)} \tag{11}$$

There we find utility attributes $x_h(a_i,\, s_j)$ and data probabilities $p(d_k|s_j)$ side by side, as factors in the same product, except that their relative impacts are dependent on w_h and $p(s_j)$, respectively. Also, $p(d_k|s_j)$ enters

the denominator whereas $x_h(a_i, s_j)$ does not. They can be considered partly compensatory for each other. It is also noteworthy that importance weights w_h for attributes in linear additive utility assessment, and probabilities $p(s_j)$ of states in expectation calculation can be treated, formally, the same way. This demonstrates how important it is to have all cues that enter the evaluation of a situation independent, no matter if they are considered "data" for probability revision, or "attributes" for multi-dimensional utility assessment, because otherwise the same evidence will enter twice into the calculation of the same EV.

Assume, for demonstration purposes, that a certain probability $p(s_j)$ and a corresponding weight on an attribute w_k are equal - e.g., this could be the case, in our used car example, if our overall objective was to obtain a certain fixed retail price. A particular state s_j among other states would, then, be the event of actually obtaining this retail price, and $p(s_j)$ or $p(d_k|s_j)$ would be the probability of this event, and $p(d_k|s_j)$ any data likelihood operating on this event probability. Similarly, w_h would be (among other aspects) the weighting factor on the attribute of retail price. A relevant datum d_k could be the information that a friend just sold a comparable car at the desired price. Now, we could use this information in assessing a quantity x_h to the value of the attribute "retail price" of the car in question, or we could use this information in modifying our probability $p(s_j)$ of obtaining the desired retail price into $p(s_j|d_k)$, but not both. If we have the choice of

using this information either way, and want to come up with the same final overall expected value for the car, it is quite clear that there is a well defined function relating $p(d_k|s_j)$ and x_h to each other. However, figuring out this relation requires a set of other restrictions and assumptions, so we will not go into more detail here.

We do have some experience about what happens with conditionally non-independent data (e.g., Domas & Peterson 1972), but we still have to find out what happens (both to the normative model and to the actual behavior of human decision makers) if the non-independence exists between probability and utility assessments, $p(d_k|s_j)$ and $x_h(a_i, s_j)$, or between the x_h's, or between probabilities and weight factors w_h. In most real life situations, many of these variables will be mutually interdependent. To give a simple example, assume you plan to go on a mountain hike, and evaluate the food you might take along. In this case, the weight factor w_h of the utility dimension "edibility without cooking" will be strongly dependent on the probability of finding conditions to set up a fire, and the weight factor w_h for (physical) weight will be dependent on the probability of finding some means of transportation for part of the way. Of course, these problems are partly related to those arising from non-independence between actions and states obtainable which have been dealt with in the paper by Luce & Krantz (1971), but partly they are of a different nature, as in the example above, in which case the normative model would have to be changed in such a way that the weight factors w_h are no longer independent of the states s_j

obtained.   As far as such dependencies can be covered by a simple additional

utility of the state s_j as such, independent of the action chosen (as in Luce

& Krantz 1971, p. 262), they are taken care of by the regret transformation.

Criticism of the assumption of continuous functions

The general model described above assumes continuous probabilities and utility cues (attributes) which are aggregated together to give DM an overall continuous evaluation $\underline{EV}(a_i)$ of his choices a_i. Small changes in $p(s_j|d_k)$, w_h or $x(a_i, s_j)$ cause small changes in $\underline{EV}(a_i)$, and DM takes these changes, or differences between $\underline{EV}(a_i)$ for different a_i into account to pick the one for which $\underline{EV}(a_i)$ is maximal over a_i.

Experimental results, however, cast some doubt on whether human $\underline{Ss}$ can do such fine-grain analyses in the first place, and if they need to in the second place. The first would be a technical problem of providing adequate decision aids, the second alludes to the problems of flat maxima, satisficing, and focussing phenomena. The multi-attribute utility model of Eq. [8] assumes linearity, additivity, and independence of the attributes h. It does not take into account interactions between variables, nor hierarchical structures that may obtain. Hierarchical utility models, and Tversky's (1971) Elimination-by-Aspects model might do a better job than the weighted-average model of Eq. [8] when it comes to these phenomena. The relationships between such lexicographic (hierarchical) and additive models has been shown by Fishburn (1970, p. 48). Another attempt to by-pass these difficulties without leaving the general framework of model [9] would be to assume non-linear non-continuous utility functions for the attributes h which allow for mutual dependencies, replacing the products $w_h\, x_h\, (a_i, s_j)$ by more

general functions the tape

$$
\phi_h(x_h(a_i, s_j)) = \begin{cases} {}_1\phi_h \ (x_h(a_i, s_j)) \ \text{if some specified conditions hold,} \\ \\ {}_2\phi_h \ (x_h(a_i, s_j)) \ \text{if some other specified condition hold,} \end{cases}
$$

[12]

when the functions ${}_e\phi_h(x_h(a_i, s_j))$ may or may not take into account the states of other attributes than x_h (thus taking care of interactions, e.g., if these other attributes enter the function multiplicatively), and where the conditions for choice of function on the right hand side may also contain other attributes.

Also, the ${}_e\phi_h(x_h)$ may be constant functions over certain intervals of x_h -- thus indicating that you don't have to care about the actual value of x_h as long as it is within a certain range which satisfies your needs.

To make these ideas a little more clear by means of an example, let's assume that you want to evaluate cars (a_i). One dimension or aspect in evaluating the options would be the attribute (x_1) of having or not having an air conditioner, and maybe the capacity or power of the air conditioner if you want to measure it continuously. Another attribute (x_2) would be convertible or not, or a sun roof. In this case, you would evaluate x_1 differently dependent on the respective states of x_2, and vice versa.

Moreover, (and this is to come back to Section III of this paper), you could eliminate these aspects completely from your utility assessment, and rather consider the different degrees of being comfortable in your car as possible states of nature (s_j), and include the facts of air conditioners, convertibles and sun roofs as data (d_k) to modify your probabilities of these possible states.

It has frequently been argued that linear additive functions ([9]) provide satisfactory approximations to the more complicated models. (Yntema & Torgerson 1961, Slovic & Lichtenstein 1971, Fischer 1972a, and Dawes 1972). However, this is true only in some special cases, and not in general. Moreover, this claim has been made mainly for the riskless case, (Fischer 1972b), whereas we are here concerned with the risky situations. For these, however, Fischer (1972b, p. 13) has pointed out that it is not so likely that the conditions for linear additive utility models hold under risky conditions. As Fischer 1972b has shown in a sensitivity analysis, the Yntema and Torgerson 1961 findings have been overinterpreted: additive approximations accounted only for 23% of the variance of composition rules involving all main effects and cross products of 9 attributes.

Deterministic vs. probabilistic models

Since the early days of phychophysics more than a century ago, psychologists have known that men do not react consistently (in a deterministic sense) to stimuli, but rather on a probabilistic basis: knowledge of stimulus conditions does not enable us to make deterministic predictions of behavior but only give us probabilities for possible actions. Thurstone (1927) and Luce (1959) have generalized these probabilistic choice models beyond psychophysics into fields of psychology where no physical measures of stimuli are available. Since then, probabilistic choice models (preference models) have been generally accepted for decision making in riskless choice situations. The concept of absolute consistency in $\underline{S}$'s choices has been replaced by a relative one, assuming stochastic transitivities rather than absolute transitivities.

But in experimental risky decision making situations, stronger criteria are applied to the data: the classical expectation maximization models tested in most decision making experiments are deterministic: in principle, they do not allow for any random fluctuation of choices, or errors. For this reason, reports on experimental results obtained in the context of such deterministic models in many cases look rather helpless, like "14 out of 20 $\underline{S}$ s displayed less than 5 of 30 possible errors" - helpless looking in the sense that the experimenter actually does not know how to evaluate his (deterministic) model in the light of these data which partly support his model but otherwise contain "errors" which the model does not allow.

This becomes a particular problem when you want to do a Bayesian data analysis - and most experimenters working in decision making are inclined to do Bayesian data analyses rather than classical statistics.

Bayesian data analysis consists, essentially, in calculating likelihood ratios

$$L = \frac{P(D|H_a)}{P(D|H_b)}$$

[13]

to test two competing hypotheses against each other, where $P(D|H_a)$ and $P(D|H_b)$ are probabilities of occurrence of data D assuming that hypothesis H_a or H_b holds, respectively. The data D consist, in most cases, of a sequence of observed choices of alternatives by the DM, or $\underline{S}$. Deterministic choice models, like the expectation maximization model, imply data probabilities (choice probabilities) of $P(D_i|H_j) = 1$ if the choice is considered optimal by the model, and 0 otherwise.

It can easily be seen that such deterministic hypotheses cannot be handled by Bayesian data analyses, since the calculations of likelihood ratios cannot be done with probabilities of 0 or 1. What is called for is some kind of error theory.

An easy way of doing this is to replace the deterministic maximization model by a probabilistic preference model, as it has been done for riskless choices by the Thurstone or BTL model. E.g., an expectation preference model (rather than expectation maximization model) could assume choice probabilities

$p(a_k)$ for alternatives (gambles) a_k

$$p(a_k) = \frac{EV_i}{\sum\limits_{i=1}^{m} EV_i}$$

[14]

where the summation of EV_i is overall alternatives offered for choice at the same presentation.

This model leads into problems if there exist choice alternatives with negative expectations. Arbitrary addition of a constant to all EV_i could make them all positive but is not permitted in the framework of the BTL model.

Transforming the payoffs $v(a_i, s_j)$ into regrets $r(a_i, s_j)$ would give all outcomes the same sign but at the same time switch the scores from desirabilities to undesirabilities, i.e. the higher the regret, the lower the probability of choosing that particular alternative. Three alternative models will be proposed for this purpose, which we might call regret-avoidance models:

(a) the sum-difference regret model:

$$p(\text{choice of } a_k) = \frac{\sum\limits_i ER_i - ER_k}{(n-1) \sum\limits_i ER_i}$$

[15]

(b) the maximum-difference regret model:

$$p(\text{choice of } a_k) = \frac{\max\limits_i [|ER_i|] - |ER_k|}{n \max\limits_i [|ER_i|] - \sum\limits_i |ER_i|}$$

[16]

(c) the reciprocal regret model:

$$p(\text{choice of } a_k) = \frac{1}{ER_k \cdot \sum_i \frac{1}{ER_i}}$$

[17]

For the regret matrix [3] from page 4, these models predict the following choice probabilities:

i	regrets		ER_i	p(choice) predicted from model		
				(a)	(b)	(c)
1	-110	0	-22	.404	.634	.496
2	-98	-21	-36.4	.341	.366	.309
3	0	-70	-56	.255	0	.195
$p(s_j)$	.2	.8				

Model (a) is the least sensitive to differences in ER, model (b) predicts always a choice probability of 0 for the worst alternative (which may lead into problems when there are $\underline{S}$s actually choosing this worst alternative occasionally), model (c) may run into a problem when there is a choice alternative with ER = 0.

None of these regret-avoidance models has been tested on empirical data, nor has any probabilistic expectation preference model, as far as I know.

Alternative approaches to adopt the deterministic decision making model for Bayesian data analysis, i.e. to get rid of choice probabilities of 0 and 1, have been initiated by Ward Edwards and Dennis Fryback in recent discussions (personal communications).

Ward Edwards suggested that $\underline{S}s$ might still choose the optimal gamble but not all the time. His model assumes that

$$p(\text{choose } a_k) = \beta_1 \left[EV(a_k) \right] + \beta_2 \left[VAR(a_k) \right] + (1-\beta_1-\beta_2) \left[\text{random process} \right]$$

[19]

with $0 \leqslant \beta_1 + \beta_2 \leqslant 1$. This means that $\underline{S}$ chooses a_k with a certain probability β_1 if its EV is maximal, or with a certain probability β_2 if its variance is ideal for $\underline{S}$, (in the sense of portfolio theory) or on basis of some random process (like, e.g., equal probabilities for all alternatives) with probability $(1-\beta_1-\beta_2)$.

Dennis Fryback suggested that $\underline{S}s$ might choose the optimal gamble with probability α, $0 < \alpha < 1$, and all other gambles with probabilities $(1-\alpha)/(n-1)$, i.e. the remaining probability $(1-\alpha)$ if $\underline{S}$ does not choose the optimal gamble is evenly distributed over all other choice alternatives.

In both the Edwards and the Fryback models, maximum likelihood estimates for the parameters β_1 and β_{2-1} or α, respectively, could be obtained from the data.

Preliminary results from re-analyzing available decision making data, i.e. choices among bets, have shown that these data still have higher likelihoods under the diluted deterministic models rather than under the probabilistic choice models (Ray Seghers) and also as compared to the one-peak model by Fryback. However, a satisfactory model to explain the deviations from the deterministic normative model has still to be developed.

References

Bernoulli, D. Specimen theorie novae de mensura sortis. Commentarii
Academiae Scientiarum Imperiales Petropolitanae, 1738 5, 175-192. Translated
in: Econometrica, 1954, 22, 23-36.

Coombs, C. H., Dawes, R. M., & Tversky, A. Mathematical Psychology: an
Elementary Introduction. Englewood-Cliffs, N.J.: Prentice Hall, 1970.

Coombs, C. H., & Pruitt, D. G. Components of risk in decision making:
probability and variance preferences. J. Exp. Psychol., 1960, 60, 265-277.

Dawes, R. H. Slitting the decision maker's throat with Occam's razor -
The superiority of random linear models to real judges. Paper presented at
Chicago School of Business, Nov. 17, 1972.

De Groot, M. H. Optimal statistical decisions. New York: McGraw-Hill, 1970.

Domas, P. A., & Peterson, C. R. Probabilistic information processing systems:
evaluation with conditionally dependent data. Organ. Beh. Hum. Per., 1972, 7,
77-85.

Dubins, L. E., & Savage, L. J. How to Gamble if You Must. New York:
McGraw-Hill.

Edwards. W. The Theory of Decision Making. Psychol. Bull. 1954, 51, 380-417.

Feller, W. An Introduction of Probability Theory and Its Applications.
New York: Wiley, 1948.

Fischer, G. W. Multi-dimensional value assessment for decision making.
Engineering Psychology Laboratory Technical Report # 037230-2-T, University of
Michigan, June 1972a.

Fischer, G. W. Four methods for assessing multi-attribute utilities: An experimental validation. Engineering Psychology Laboratory Technical Report #037230-6-T, University of Michigan, Sept. 1972b.

Fishburn, P. C. Utility Theory for Decision Making. New-York:Wiley, 1970.

Fryback, D. G., Goodman, B. C., & Edwards, W. Choices among bets by Las Vegas gamblers: Absolute and contextual effects. J. Exp. Psychol., 1973, in press.

Luce, R. D. Individual Choice Behaviors. New York:Wiley, 1959.

Luce, R. D., & Krantz, D. H. Conditional expected utility. Econometrica, 1971, 39, 253-271.

Pollatsek, A., & Tversky, A. A theory of risk. J. Math. Psychol., 1970, 7 (3), 540-553.

Restle, F. Psychology of Judgement and Choice. New York:Wiley, 1961.

Savage, L. J. The Foundations of Statistics. New York:Wiley, 1954.

Schneeweiß, H. Entscheidungskriterien bei Risiko.(Ökonometrie und Unternehmensforschung VI). Springer-Verlag, Berlin-Heidelberg-New York, 1966.

Slovic, P., & Lichtenstein, S. Comparison of Bayesian and regression approaches to the study of information processing in judgment. Oregon Research Institute Research Monograph, Vol. 10, No. 1. Organ. Beh. Hum. Per., 1971, 6, 649-744.

Thurstone, L. L. A law of comparative judgment. Psychol. Rev., 1927, 34, 273-286.

Tversky, A. Elimination by aspects: A probabilistic theory of choice. Michigan Mathematical Psychology Program MMPP 71-12, 1971.

von Neumann, J., & Morgenstern, O. Theory of Games and Economic Behavior. Princeton:University Press, 1944.

von Winterfeldt, D., & Edwards, W. Costs and payoffs in perceptual research. Engineering Psychology Laboratory, Unpublished paper, University of Michigan, 1972.

Wendt, D. Utility and risk. Acta Psychol., 1970, 34, 214-228.

Winkler, R. L., & Murphy, A. H. Experiments in the laboratory and the real world. Organ. Beh. Hum. Per., 1973, in press.

Yntema, B. D., & Torgerson, W. S. Man-Computer Cooperation in Decisions Requiring Common Sense. IRE Transactions on Human Factors in Electronics, 1961, HFE-2, 20-25.