Constructive Decision Theory

Glenn Shafer
Department of Mathematics
University of Kansas
Lawrence, Kansas   66045
December 1982

Working paper.   Comments welcomed.

Constructive Decision Theory

A fully constructive decision theory should show us how to make the probability and value judgments on which a given decision can be based. It should not rule out the possibility that some of these judgments have already been made--that we have already formulated some relevant beliefs and preferences. But it must also deal with the possibility that all the needed probabilities and values must be constructed ab initio, by weighing factual and moral arguments.

Can Bayesian decision theory or subjective expected utility theory, as it is sometimes called, be success-fully interpreted as a fully constructive theory? Or does the constructive idea lead us to alternative decision theories?

In this paper, I discuss the difficulties in inter-preting expected utility theory constructively and argue that these difficulties are best resolved by shifting to a vocabulary in which we speak of values for goals instead of utilities for consequences. This leads to a theory of decision that meshes with the theory of probability judg-ment developed in my book, A Mathematical Theory of Evidence (1976).

1. <u>Bayesian Decision Theory</u>:   <u>Descriptive, Normative</u>,
   <u>or Constructive</u>?

   Advocates of Bayesian decision theory generally
follow L. J. Savage (<u>The Foundations of Statistics</u>,
1954) in emphasizing that the theory's axioms or
postulates are normative, not descriptive. They
agree with Savage that though these axioms can be
interpreted as a "crude and shallow empirical theory,"
(p. 20), they are better interpreted as a norm, a
standard of coherence and consistency that thought-
ful decision makers should strive to attain.

   Unfortunately, the words "normative," "coherence,"
and "consistency" have been the source of much misunder-
standing and ill will. They give the impression that
Bayesians see their axioms as the only sensible way
to approach decision making. This impression is some-
times justified; many Bayesians do believe that there
is no sensible alternative to their approach and do use
these words as epithets--insults directed towards statis-
ticians who use other approaches to probability judg-
ment and decision. But some Bayesians think of their
axioms as "normative" only in the sense of recognizing
that trying to construct preferences satisfying them
is one way to make thoughtful deicsions. By calling

the theory "constructive" rather than "normative", we emphasize this latter, undogmatic approach to the axioms.

From a constructive point of view, Bayesian decision theory breaks a decision problem down by decoupling our consideration of probabilities from our consideration of utilities. Probabilities and utilities are constructed separately and then combined in the calculation of expected utilities for various possible decisions (cf. Raiffa, pp. 127-128.) But how do we construct probabilities and utilities?

Constructive Probability. How do we construct probabilities? Much of the work done on this question has used the vocabulary of elicitation. (See, for example, Savage 1971, and Kadane, et. al. 1980.) We assume, tongue-in-cheek, that we already have probabilities in the back of our minds satisfying the Bayesian axioms, and we ask ourselves questions calculated to reveal or elicit these probabilities.

Elicitation can produce answers, but it is not clear how to get our tongue back out of our cheek. What meaning and what persua-sasiveness do the answers have once it is admitted that there really are no pre-determined probabilities in the back of our minds? This question becomes especially embarrassing when methods for reconciling inconsistent elicitations are based on full-blown statistical models for "errors of elicitation"--models, that is to say, for how elicited probabilities differ from the fictitious true probabilities in the backs of our minds. (Models for errors in elicitation have been

proposed by Lindley, Tversky and Brown 1979, and by Dickey 1980. Both of these papers are followed by a discussion that wrestles with the fictitious nature of these errors.)

We need some new ideas about how to construct probability distributions. Most of all, we need an explicitly constructive framework for these ideas. We need a vocabulary that allows us to talk forthrightly about constructing probability judgments.

Shafer and Tversky (1983) attempt to provide such a constructive framework and vocabulary. Their basic idea is that Bayesian probability judgment consists of matching the evidence in particular problems to a scale of canonical examples where the truth is determined by chance. The matching process involves interaction with the evidence--we sample our evidence, examine it from different viewpoints and search for ways to make it fit on the scale. Moreover, the matching process is always based on a design--a specification of how a whole probability distribution is to be built up from elementary probability judgments.

This constructive framework has two immediate virtues. First it makes clear the possibility of alternative theories. The Bayesian theory corresponds to one scale of canonical examples, but there are other, perhaps equally useful scales. Secondly, it leaves us at peace with the idea that some probability judgments are better than others. Sometimes the fit to the scale is clear and convincing, sometimes not.

Can Utilities be Constructed?   The Problem of Small Worlds.
In order to use Bayesian decision theory, we need to construct not
only probabilities but also utilities. Given a list of possible
ways things might turn out, we need to assign each a numerical
utility.

But here we encounter a fundamental difficulty. It seems
impossible to assign a utility to any particular description of how
things might turn out without implicitly assessing probabilities
for how yet further matters will turn out.

Savage gives the following example:

...Jones is faced with the decision whether to buy a certain
sedan for a thousand dollars, a certain convertible also for
a thousand dollars, or to buy neither and continue carless.
The simplest analysis, and the one generally assumed, is that
Jones is deciding between three definite and sure enjoyments,
that of the sedan, the convertible, or the thousand dollars.
Chance and uncertainty are considered to have nothing to do
with the situation. This simple analysis may well be appro-
priate in some contexts; however, it is not difficult to
recognize that Jones must in fact take account of many uncer-
tian future possibilities in actually making his choice. The
relative fragility of the convertible will be compensated only
if Jones's hope to arrange a long vacation in a warm and scenic
part of the country actually materializes; Jones would not buy

a car at all if he thought it likely that he would immediately
be faced by a financial emergency arising out of the sickness
of himself or of some member of his family; he would be glad to
put the money into a car, or almost any durable goods, if he
feared extensive inflation. This brings out the fact that what
are often thought of as consequences (that is, sure experiences
of the deciding person) in isolated decision situations typically
are in reality highly uncertain. Indeed, in the final analysis,
a consequence is an idealization that can perhaps never be well
approximated.... (Foundations, pp. 83-84).

One might think, at first glance, that a sufficiently detailed
description of Jones's possible future situations would make it
possible for him to decouple his utilities from his probabilities.
But in practice it seems impossible to give sufficiently detailed
descriptions. No matter how much detail we include in a description
of a situation, there always remain uncertainties that can affect the
degree to which we will enjoy or value that situation. This is the
point of the last sentence in the passage above. Savage used the
term "consequence" to refer to a description to which a person can
attach a utility independently of any residual uncertainties. But he
had to admit that this is an idealization "that can perhaps never be
well approximated."

Savage used the term "small world" to refer to a set of possibili-
ties that are described in limited detail. So we may call the problem

we have just described "the problem of small worlds": pure judg-
ments of utility seem impossible because closer examination always
shows the utility of a thing or a situation to depend on further
judgments of probability.

The problem of small worlds will not trouble us so long as
we use the vocabulary of elicitation. If we pretend that preferences
as detailed as one pleases are already determined and waiting to be
elicited from the backs of our minds, then we will see no reason to
be disturbed that particular utilities we elicit are, from a deeper
point of view, expected utilities. But there does seem to be a pro-
blem when we use the vocabulary of construction.

It is tempting to describe the problem of small worlds as a
problem of consistency. How do we know that the utility we construct
at one level of description will suit us when we move to a more
detailed level of description? How do we know, that is to say, that
it will agree with the expected utility that we calculate after
constructing probabilities and utilities at this more detailed level?
But it may be more enlightening to say that the problem is that we are
unable to construct a utility in the first place. The language of
utility suggests that we are supposed to ask ourselves how much a
given thing or situation is worth to us--how much we will like it.
The problem is that we can never give a straight answer to this question.
We always have to say "that depends" and then turn to constructing
probabilities for the events it depends on.

Savage on Small Worlds

Since I have borrowed the term "small world" from Savage, I should point out that the problem of small worlds I have just described is quite different from the problems Savage dealt with in his discussion of small worlds on pages 82-91 of Foundations.

The difference between Savage's problem of small worlds and ours is one result of a fundamental difference between Savage's viewpoint and our constructive viewpoint. Savage begins with the elicitation of preferences. He assumes the existence of definite preferences between acts and then formulates postulates governing these preferences that allow him to derive probabilities and utilities from them. Our viewpoint directs attention in exactly the opposite direction. Our thought is to first construct probabilities and utilities and then possibly use these to construct preferences between acts.

Savage's problem of small worlds arises because given preferences between acts, when analyzed in terms of two different small worlds, one larger (i.e., more detailed) than the other, may yield inconsistent probabilities. Since the probabilities derived from the smaller world then appear erroneous, Savage's problem is to decide when a small world is detailed enough to give correct probabilities. (Foundations, pp. 89-90.)

Savage's viewpoint is discussed more fully in an appendix below.

2. <u>Utilities for Consequence vs. Values for Goals</u>

The reader has probably reacted to my contention that the problem of small worlds prevents a person from ever making judgments of utility by pointing to the fact that people do make judgments of utility. Decision analysts do assess utilities as well as probabilities. And even people untutored in decision theory manage to decide what they want.

I would like to suggest that when people do construct utilities they avoid the problem of small worlds simply by deciding to value the "consequence" to which they are attaching the utility irrespective of whatever else happens.

"I have decided to buy a convertible," Jones tells us, "because my wife and I are taking a vacation to New Mexico this summer, and we really want to enjoy the sun." "You should think this through more carefully, Jones," we respond. "Don't you remember that sunburn you got on the beach at Fort Lauderdale last summer? You never really enjoy these vacations anyway. And if your wife does like the sun that much, she may not come back to Chicago with you." "You are always dreaming up things to worry about," replies Jones. "I detest this winter weather, and I have set my heart on a tour of the desert in the sun. The trip may be a disaster, but staying home might be a disaster, too. Who knows?"

Jones has decided on a trip to sunny New Mexico in a convertible. He does not want to analyze all the different ways taking the trip or

not taking the trip might turn out, partly because he does not feel he can construct convincing probabilities for them, but also because these more detailed scenarios are not really the object of his desire. The trip lies within the bounds of accepted behavior, and he and his wife have decided they want to go.

I would like to suggest that we can better discuss Jones's behavior, and better order our own, if we abandon the vocabulary of utility and talk instead about goals. Jones and his wife have made a trip in the sun their goal. Rational decision in general, I would argue, begins with the construction of goals.

The vocabulary of goals has several advantages over the vocabulary of utility for a constructive theory of decision. First, it keeps the need for construction always before our eyes; goals obviously have to be made. Second, it fits the idea of working at a given level of description and therefore avoids utility's problem of small worlds. Third, it steers us well clear of the hedonistic interpretation into which talk of utility always threatens to slip.

Constructivity. The idea of utility has its historical roots in the descriptive science of economics: it represents an effort to find an explanation of human action that lies deeper than superficial goals and motives. It is difficult, therefore, to sustain a constructive attitude when we use the vocabulary of utility. As we try to construct at a given level, the vocabulary is suggesting that there somehow already exists a better-structured reality at a deeper level.

The vocabulary of goals can free us from this picture. It can
allow us to see that the goals produced by conscious thought
and deliberation are the clearest and most definitely structured
of all our goals and motives. What lies deeper is not well-structured
hidden utilities, but rather the raw material from which we have
built. This raw material can include impulses and passions, ideas,
good and bad examples, and biological and physical constraints.

Limited Description. The idea of utility constantly pushes
us toward more detailed descriptions: we can assign a pure utility
to a situation only if the situation is fully described. But goals
are more limited and concrete. We can describe succinctly what it
means to take a sunny vacation in New Mexico, to win a certain game,
to get a certain job, to raise a good crop, or to save a child from
a life-threatening disease.

Hedonism. The theory of subjective expected utility is not
officially hedonistic. And when the theory is interpreted descrip-
tively, it is clear that utility has nothing to do with pleasure;
utility is simply a mathematical construct that results from the
axioms on preferences being satisfied. But when we try to inter-
pret the theory constructively, a hedonistic conception of utility
often seems to lurk near the surface. When we are given a detailed
description of how things might turn out and are asked to say how
much utility we assign to that state of affairs, how can we answer
except by trying to imagine how much we would enjoy it? When we
think in terms of goals, on the other hand, we are not so likely to

be forced into a hedonistic interpretation.

The achievement of goals does often involve pleasure, of course. But it is a mistake to suppose that a goal is only the surface manifestation of an underlying pursuit of pleasure. Often the achievement of a goal is pleasurable only as a consequence of our having made it a goal: running five miles without stopping can bring delight to a person who has just taken up jogging, whereas it would have only been painful to the same person were he forced to do it as a punishment. Moreover, we make and pursue goals fully realizing that their achievement may be spoiled by other unhappy events.

3. Goals as Subsets

What is the technical significance of attaching values not to consequences, which become ever more specific as descriptions become more detailed, but to goals, which can be fully specified at a given level of detail? In terms of a frame of discernment, the significance is that values are attached not to points but to subsets.

Frames of Discernment. Savage and most of his successors use two sets of descriptions in their decision theories: a set C of consequences and a set S of states of the world. (See Fishburn, 1981.) The consequences are "states of the person" (Foundations, p. 13): they specify how things that the person cares about turn out. The states of the world specify more general facts about the world, facts which determine what consequences follow from various acts. (See the

appendix below for a further discussion of Savage's approach.) Since
our purpose is to discuss goals rather than to formulate axioms govern-
ing preferences between acts, we will find it more convenient to deal
with a single set. Following the practice in my book A Mathematical
Theory of Evidence (1976, Chapter 2), we may call this set a "frame
of discernment," or simply a "frame."

A frame is a set of descriptions of how things turn out. These
descriptions are, in the opinion of the person who constructs or
formulates the frame, mutually exclusive and collectively exhaustive;
the person knows that exactly one of the descriptions will turn out
to be correct.

It is often necessary to refine a frame. This means extending
each of its descriptions--i.e., making each description more detailed.
We usually have to consider several incompatible ways in which a
description can be extended, so in general refining a frame means
replacing each description by a set of descriptions. We say that a
frame Ω is a "refinement" of a frame Θ (and that Θ is a "coarsen-
ing of Ω) if each description θ in Θ is consistent with a
non-empty subset, say $\omega(\theta)$, of descriptions in Ω , and the sets
$\omega(\theta)$ form a disjoint partition of Ω .

A subset A of a frame Θ can be thought of as a proposition--
the proposition that one of the descriptions in A is true. If Ω
is a refinement of Θ , then a subset A of Θ corresponds to the
same proposition as the subset $\underset{\theta \in A}{\cup} \omega(\theta)$ of Ω .

<u>Goals.</u>

Goals also correspond to subsets of frames. In order to express a given goal, we need, of course, to choose a fine enough frame, a frame whose descriptions are sufficiently detailed to take the goal into account. Further refinement will then make no difference; if a given goal corresponds to the subset A of a frame Θ, then it will also correspond to the corresponding subset of Ω, namely $\bigcup_{\theta \in A} \omega(\theta)$.

Suppose we specify n goals, corresponding to n subsets of a frame Θ. Let us say that these goals are "consistent" if the subsets all overlap--i.e., their intersection is non-empty. And let us say that they are "monotonic" if the subsets are nested-- i.e., if they can be ordered and labeled $A_1, \ldots, A_n$ in such a way that $A_1 \supset A_2 \ldots \supset A_n$. Notice that consistency and monotonicity are unaffected by refinement.

When we construct goals we usually try to make them consistent. It is reasonable, however, to admit the possibility of inconsistent goals. One argument for doing so is provided by the fact that initially consistent goals may become inconsistent as a result of new knowledge. After learning more about the difficulties involved in building a barn, a person may realize that he cannot both build a barn for his cow before winter and also take a vacation to Yellowstone. This is illustrated in Figure 1. The new knowledge eliminates the possibilities in $\overline{\Theta}_0$, thus reducing Θ to Θ_0 and reducing the

subsets A_1 and A_2 to $A_1 \cap \Theta_0$ and $A_2 \cap \Theta_2$. The goals were initially consistent: $A_1 \cap A_2 \neq \emptyset$. But now they are inconsistent: $(A_1 \cap \Theta_0) \cap (A_2 \cap \Theta_0) = \emptyset$.
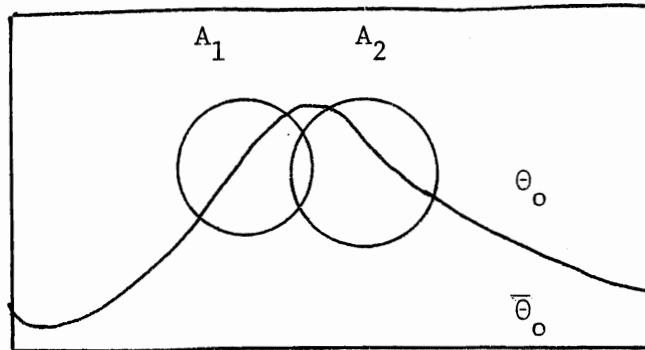

Figure 1.

The significance of monotonicity is that monotonic goals cannot be made inconsistent by new knowledge. (See Shafer 1976, p.221.) This point is illustrated by Figure 2. There are initially four nested goals in this figure: $A_1 \supset A_2 \supset A_3 \supset A_4$. Reduction to Θ_0 makes A_4 impossible and thus eliminates it as a goal, but the remaining goals are still nested: $A_1 \cap \Theta_0 \supset A_2 \cap \Theta_0 \supset A_3 \cap \Theta_0$.

It is always our privilege to change the goals we have constructed, and in particular we always can, if we wish, reformulate our goals so that they will we monotonic. Suppose, for example, that we at first
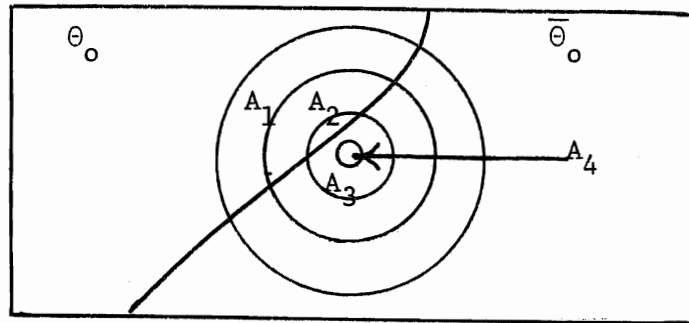

Figure 2.

construct overlapping goals A_1 and A_2 , as in Figure 3, and
that we consider goal A_1 more important than goal A_2 .
Perhaps A_1 is the goal of restoring a friendship that has been
interrupted by a misunderstanding, while A_2 is the goal of having
oysters for supper. We can, if we wish, replace these two goals
with a set of related but monotonic goals: $B_1 = A_1 \cup A_2$, $B_2 = A_1$,
and $B_3 = A_1 \cap A_2$. From a purely aesthetic point of view, this
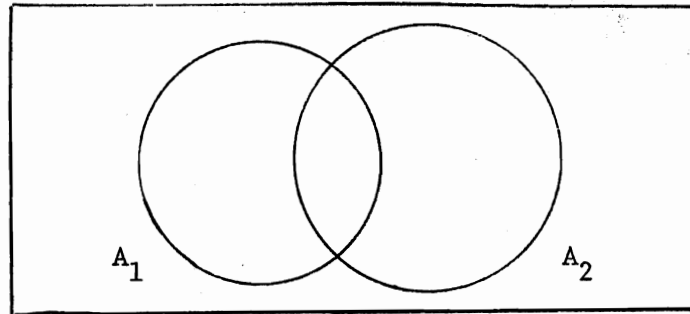replacement does not seem attractive;

Figure 3.

It is awkward to say that one of our goals is to either restore a
friendship or have fresh oysters for supper. More important than
the aesthetic point of view, however, is the point of view that sees
goals as the objects of actions. Here the question is whether we
are interested in an action that promises to either restore the friend-
ship or produce the fresh oysters. We will return to this question
shortly.

We have stressed that if we formulate goals with respect to one
frame and then refine the frame, the same goals will still appear
in the refinement. This is the advantage of goals over utilities;

The problem of small worlds is that utilities with respect to one frame always turn out to be expected utilities with respect to some refinement. But refinement does bring new issues into the discussion and therefore may lead us to formulate new goals. This potential for new goals is one reason why an insistence on monotonicity for goals does not seem desirable. The addition of a new goal to a monotonic set of goals will usually destroy the monotonicity, and so an insistence on monotonicity will mean that refinement will entail not only the construction of new goals but also the constant reformulation of already constructed goals.

Richard Jeffrey's "Logic of Decision." In abandoning Savage's distinction between consequences, and states of the world and using a single set to carry both probability and value, we are following the lead of Richard Jeffrey. In The Logic of Decision (1965); Jeffrey develops a theory of expected utility in which probabilities, utilities, and preferences are all attached to propositions, expressed as subsets of a single frame.

Jeffrey also incorporates the actions we are considering into his frame. He requires that each description in the frame specify which action we take, so that the proposition that we take a given action corresponds to a subset of the frame. We should, perhaps, follow Jeffrey's lead in this respect as well. We do sometimes want to attach moral values to certain actions without regard to what these actions achieve, and we can do so within the constructive

theory we are developing if the action appears as a proposition in the frame. In the present paper, however, I will consider only a simple set-up where actions are not incorporated into the frame.

It should be noted that Jeffrey's way of attaching value to a proposition is quite different from ours. We propose to call a proposition a goal and, treating it as a goal, to attach a value to it that is unaffected by opinions about what else is likely to be true if the proposition is true. The utility that Jeffrey attaches to a proposition, on the other hand, is best described as a person's "expected utility given the proposition is true"; it does depend on what else the person thinks is likely to be true if the proposition is true.

A number of authors who have found Savage's theory attractive have found Jeffrey's more holistic theory puzzling and unappealing. (See, for example, Fishburn 1981.) This may be partly explained by saying that Jeffrey's theory is more obviously non-constructive than Savage's. We are accustomed enough to expressing preferences between acts that the idea of using such preferences as building blocks in the construction of probabilities and utilities is at least super-ficially plausible. We are much less accustomed to expressing the kind of preference between propositions that Jeffrey requires. Moreover, when we think about going the other way, from probabilities and utilities to preferences, the problem of small worlds appears much closer to the surface in Jeffrey's theory than in Savage's.

By isolating "consequences" from "states of the world," Savage
makes it plausible that we can make pure judgments of utility. By
attaching utilities to the subsets as well as the points of our
frame, Jeffrey makes it obvious that all these utilities involve
judgments of probability.

4. Relating Actions to Goals

Suppose we have chosen goals corresponding to subsets
$A_1,\ldots,A_n$ of a frame Θ . Consider a set Δ of possible actions.
How can we describe mathematically the relation between the actions
in Δ and the goals $A_1,\ldots,A_n$?

The simplest situation is one where we know what each action
would accomplish. Suppose there corresponds to each action δ in
Δ a subset $A(\delta)$ such that performing the action δ will, we know,
assure that things will turn out according to one of the descriptions
in $A(\delta)$. In this case we will look to see which goals each of
the actions will achieve, and we will also look to see whether any
of the actions will preclude any of the goals. Each goal is an
argument for one of the actions that achieves it and an argument
against any actions that preclude it.

Counting Goals. Often the goals we are discussing all carry
about equal weight with us. In this case we can base our evaluation
of an action on counts of the number of goals it affects.

The simplest way to think about an action δ is to count the

the number of goals it achieves,

$$u^+(\delta) = u^+(\delta;\{A_i\}) = \#\{i \mid A(\delta) \subset A_i\},$$

and the number it precludes,

$$u^-(\delta) = u^-(\delta;\{A_i\}) = \#\{i \mid A(\delta) \cap A_i = \emptyset\}.$$

If we want to assign each action a single score, then we might
assign δ the score

$$U(\delta) = u^+(\delta) - u^-(\delta),$$

which can be described as the net number of goals achieved by δ .

A more subtle approach is to think about an action δ in
terms of the minimum number of achievements and maximum number of
failures it entails. If we take the action δ , then we are sure
to achieve at least

$$v_*(\delta) = v_*(\delta;\{A_i\}) = \min_{\theta \in A(\delta)} \#\{i \mid \theta \in A_i\}$$

and at most

$$v^*(\delta) = v^*(\delta;\{A_i\}) = \max_{\theta \in A(\delta)} \#\{i \mid \theta \in A_i\} .$$

of our goals. If we want to assign each action a single score, then
we might assign δ the score

$$V(\delta) = v_*(\delta) - (k - v^*(\delta)),$$

where $k = \max_{\theta \in \Theta} \#\{i \mid \theta \in A_i\}$ is the greatest number of goals that can be
achieved. (If $A_1,\ldots,A_n$ are consistent, then $k = n$).

Using the pair $(v_*, k - v^*)$ instead of $(u^+, u-)$ corresponds, in
a sense, to monotonizing the original goals $A_1,\ldots,A_n$. This can be
made precise as follows. Set

$$B_j = \{\theta \in \Theta \mid \#\{i \mid \theta \in A_i\} \geqslant j\}$$

for $j = 1, \ldots, k$. Then B_j is the "goal" that at least j of the original goals $A_1, \ldots, A_n$ be achieved, and $B_1 \supset B_2 \supset \ldots \supset B_k$.

Theorem 1. $v_*(\delta; \{A_i\}) = u^+(\delta; \{B_j\})$ and $k - v^*(\delta; \{A_i\}) = u_-(\delta \{B_j\})$.

Proof: The condition $A(\delta) \subset B_j$ is equivalent to the condition that $\#\{i \mid \theta \varepsilon A_i\} \geqslant j$ for all $\theta \varepsilon A(\delta)$, and hence also to the condition that $v_*(\delta; \{A_i\}) \geqslant j$. So

$$u^+(\delta; \{B_j\}) = \#\{j \mid A(\delta) \subset B_j\} = \#\{j \mid v_*(\delta; \{A_i\}) \geqslant j\}$$
$$= v_*(\delta; \{A_i\}).$$

Similarly, the condition $A(\delta) \cap B_j = \emptyset$ is equivalent to $v^*(\delta, \{A_i\}) < j$, so that $u^-(\delta; \{B_i\}) = \#\{j \mid A(\delta) \cap B_j = \emptyset\} = \#\{j \mid v^*(\delta; \{A_i\}) < j\}$

$$= k - v^*(\delta; \{A_i\}). \quad \blacksquare$$

Is it more sensible to evaluate actions using scores such as (u^+, u^-), or should we instead use scores such as (v_*, v^*)? Should we, in other words, monotonize our goals? The issue is illustrated by Figure 4. Here we have two equally valued goals A_1 and A_2--say the two goals are having fresh oysters for supper and spending the evening playing with one's children. And we have two actions δ and δ'. The action δ assures one of the fresh oysters without precluding the evening with the children, while the action δ' assures only that at least one of the two goals will be achieved. If we use the scores (u^+, u^-), then δ scores better than δ', because δ achieves A_1, while δ' achieves neither goal. ($u^+(\delta) = 1$, $u^+(\delta') = 0$, $u^-(\delta) = u^-(\delta') = 0$.) But if we monotonize, or take $B_1 = A_1 \cup A_2$ and $B_2 = A_1 \cap A_2$ as our goals, then δ and δ' score equally well, because

they both achieve the goal B_1. ($v_*(\delta) = v_*(\delta') = 1$, $v^*(\delta) = v^*(\delta') = 2$.) Which approach is more sensible?

Many readers will lean towards monotonization. They will argue that though the δ is more satisfying than δ' in that it does definitely achieve one of the original goals, δ' should be given credit
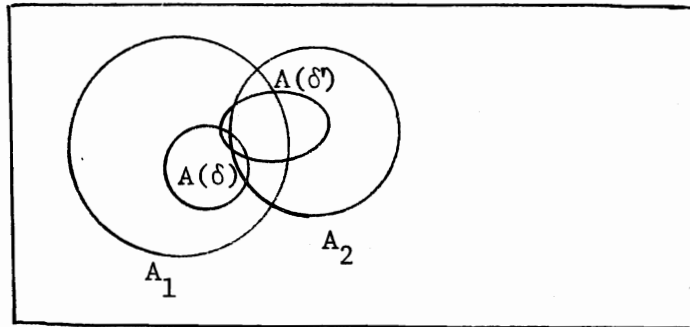

Figure 4.

for achieving $B_1 = A_1 \cup A_2$. If we score using (u^+, u^-), δ' is given no credit at all.

I would like to question this line of thought. I would like to argue that our inclination to give δ' credit for achieving $A_1 \cup A_2$ is due to the influence of a hedonistic idea of utility. If A_1 and A_2 are valued merely because they each produce one utile of pleasure for us, then achieving $A_1 \cup A_2$ is just as good as achieving A_1--either way we are assured of one utile. But if we move away from utility and think in terms of goals, it is not so clear that credit should necessarily be given for achieving $A_1 \cup A_2$. I would like to suggest that we should give credit for achieving $A_1 \cup A_2$ only if we explicitly adopt $A_1 \cup A_2$ as a goal, and that the adoption of $A_1 \cup A_2$ as a goal is

not implied by the adoption of A_1 and A_2 as goals.

People enjoy playing with their children, but they know that the value of doing so is not measured solely by the pleasure (or pain) involved. Part of the value lies in carrying out one's role as a parent. And this aspect of value involves the connection between action and goal. We value an action because of its deliberate relation to a goal. The decision to spend the evening with one's children earns directly the value we attach to the goal. Ending up spending the evening with one's children because of some complicated gambit involving oysters may not have quite the same merit.

To summarize: The scores (u^+, u^-) seem more appropriate than the scores (v_*, v^*) because they relate to actual goals. It is legitimate to monotonize or otherwise reformulate goals, but monotonization should be thought of as a deliberate and optional act, not to be taken for granted.

Weighting Goals. When we count goals we are giving the goals equal weight. A natural and hence very old idea is to consider assigning unequal weights to goals. Dawes and Corrigan (1974, p. 95) quote the following passage in a 1772 letter from Benjamin Franklin to Joseph Priestly (published by Bigelow, 1887, p. 522):

> I cannot, for want of sufficient premises, advise you *what*
> to determine, but if you please I will tell you *how* . . . .
> My way is to divide half a sheet of paper by a line into two
> columns; writing over the one *Pro*, and over the other *Con*.
> Then, doing three or four days' consideration, I put down under

the different heads short hints of the different motives,
that at different times occur to me *for* or *against* the
measure. When I have thus got them all together in one
view, I endeavor to estimate the respective weights . . . .
[to] find at length where the balance lies . . . . And,
though the weight of reasons cannot be taken with the precision
of algebraic quantities, yet, when each is thus considered,
separately and comparatively, and the whole matter lies
before me, I think I can judge better, and am less liable
to make a rash step: and in fact I have found great
advantage for this kind of equation, in what may be called
moral or prudential algebra. . . .

Suppose we attach weights $w_1, \ldots, w_n$ to the goals $A_1, \ldots, A_n$,
indicating the relative importance of these goals. The total weight
of the goals achieved by the action δ is then

$$u^+(\delta) = \Sigma\{w_i \,|\, A(\delta) \subset A_i\},$$

and the total weight of the goals precluded is

$$u^-(\delta) = \Sigma\{w_i \,|\, A(\delta) \cap A_i = \emptyset\}.$$

If we take the action, then we are sure to achieve goals of total
weight at least

$$v_*(\delta) = \min_{\theta \in A(\delta)} \Sigma\{w_i \,|\, \theta \varepsilon A_i\}$$

and at most

$$v^*(\delta) = \max_{\theta \in A(\delta)} \Sigma\{w_i \,|\, \theta \varepsilon A_i\}.$$

The use of the scores (v_*, v^*) now corresponds to the monotonization that replaces $A_1, \ldots, A_n$ with

$$B_j = \{\theta \,|\, \Sigma\{w_i \,|\, \theta \varepsilon A_i\} \geq t_j\},$$

for $j = 1, \ldots, k$, where $t_1 < t_2 < \ldots < t_k$ are the distinct values that can be obtained for the sum $\Sigma\{w_i \,|\, \theta \varepsilon A_i\}$ by varying θ.

The Case for Counting. Should we try to weight arguments, or should we simply count them, thus giving them equal weight? Several authors have argued that it is often wisest to settle for equal weights. Two arguments have been advanced.

(1) Equal weights are natural to use when we cannot decide or agree which goals to weight the most.

(2) It is natural for us to formulate goals that strike us as roughly comparable; goals we consider relatively unimportant we are likely to set aside altogether, and goals we consider especially important we are likely to analyze into a number of goals.

Argument (1) is most immediately convincing in problems of social decision, where different people disagree on what goals to weight most heavily. (See, for example, the discussion by Hammond and Adelman 1976, cited by Dawes 1979, on the choice of bullets by the Denver police. Since participants in the decision disagreed on the relative importance of preventing the person shot from returning fire, minimizing injury to the person shot, and preventing harm to bystanders, these three goals were weighted equally.) But from the constructive

point of view, individual decisions are not so different from social decisions. Individual human beings often cannot settle in their own minds the relative importance of competing goals and hence may resort to giving them equal weight.

Closely related to the problem of choosing weights for goals is the problem of choosing weights when averaging predictor variables. Here also there is an interesting literature that argues for equal weights--partly on the ground that the evidence is usually inadequate to support a choice of weights and partly on the ground that the choice is unlikely to make much difference. (See Dawes 1979 and the references therein, especially Wainer 1976.)

An important and convincing example of people's tendency to count goals is provided by recent studies of the decision-making of voters in United States presidential elections. (See Brody and Page 1973, Kelley and Mirer 1974, Shaffer 1972, and especially, Kelley 1983.) For the past thirty years interviewers from the Center for Political Studies of the University of Michigan have included in their pre-election interviews requests that voters list their likes and dislikes of the major parties and the good and bad points of the major candidates. Studies of the results indicate that a voter's eventual choice is best predicted by a scoring system that assigns to each candidate the total net number of favorable responses. Follow-up interviews have shown that these scores predict the voters' choices even better than their own statements of their intentions.

This example fits our formalism in a reasonably natural way. The good and bad points of a candidate generally correspond to the voter's goals: a president with integrity, a president who will work for higher price supports, etc. Hence a candidate's number of good points is u^+, the number of goals the voter thinks will be achieved by election of the candidate, and his number of bad points is u^-, the number of goals his election will preclude. Notice that a given goal affects a candidate's score, $U = u^+ - u^-$, only when the voter has an opinion about whether election of that candidate will achieve or preclude the goal. If the voter knows one candidate favors higher price supports but is uncertain about a second candidate's stand on the issue, he may list the issue as one of the first candidate's good points while omiting it altogether from his list of the second candidates good and bad points.

For Want of Evidence. In general, a given goal A_i does not enter into the calculation of the score $U(\delta)$ for an action δ unless the person calculating the score knows that δ will achieve or preclude the goal--i.e., unless $A(\delta) \subset A_i$ or $A(\delta) \cap A_i = \emptyset$. This feature of our calculus obviously does not accord with Bayesian ideas. A Bayesian would demand that the person supplement his opinion that δ will have the effect $A(\delta)$ with a subjective probability distribution for just which element of $A(\delta)$ will turn out to be the truth if the action δ is taken. This would lead to probabilities for whether each goal A_i would be achieved or not if the action δ were taken. The Bayesian would feel that ignoring

the goal A_i is appropriate only if this probability happens to equal $\frac{1}{2}$.

Our attitude is that though there may be evidence to support probability judgments of the kind demanded by the Bayesian, this is not necessarily the case. It is entirely possible that a voter may have no evidence whatsoever about a given candidate's stand on agricultural price supports, and in this case there will be no way to construct Bayesian probability judgments on the question.

Some readers may object to our attitude on the grounds that there surely are probabilities and that some effort should be made to take them into account. The constructive view suggests, however, that there are no subjective probabilities for an event unless and until probability judgments are made.

5. Belief Functions

We can use belief functions to make probability judgments based on limited evidence about what will happen if an action is taken.

Suppose indeed, that there are several different possibilities for what the effect of the action will be, and suppose we assess the evidence for these possibilities by assigning probabilities $p_1, \ldots, p_r$ to subsets $B_1, \ldots, B_r$; we assign probability p_j to the effect of the action being to assure the truth of B_j. This determines a "belief function" on Ω --a function Bel that assigns degree of belief

$$Bel(A) = \Sigma\{p_j | B_j \subset A\}$$

to A's turning out to be true if the action is taken. (See Shafer 1976, Chapter 2.) If $r = 1$, $p_1=1$, and $B_1=A(\delta)$, then we are back to the case of the preceding section, where we are sure only that the action will make $A(\delta)$ true. In this case we have $Bel(A(\delta)) = 1$, but $Bel(A) = 0$ for any proper subset A of $A(\delta)$.

From a constructive point of view, the subjective probability or degree of belief $Bel(A)$ given by a belief function Bel has a very different meaning from a Bayesian subjective probability $P(A)$. The Bayesian probability can be thought of as the result of matching one's evidence to a scale of canonical examples where the truth is generated by known chances: one's evidence about the action is being compared to knowledge that the action has a chance $P(A)$ of making A happen. The degree of belief $Bel(A)$ can be thought of as the result of matching one's evidence to a scale of canonical examples where the meaning of a message depends on chance: one's evidence is being compared to a message that has chance p_j of meaning that B_j will result if the action is taken. (See Shafer 1981 or Shafer and Tversky 1983.)

Suppose, as before, that we have goals $A_1,\ldots,A_n$ with weights $w_1,\ldots,w_n$. Then the expected total weight of goals achieved by the action δ is

$$E(u^+(\delta)) = \sum_{j=1}^{r} p_j \, \Sigma\{w_i | B_j \subset A_i\}, \qquad (1)$$

and the expected total weight of goals precluded is

$$E(u^-(\delta)) = \sum_{j=1}^{r} p_j \, \Sigma\{w_i | B_j \cap A_i = \emptyset\}. \qquad (2)$$

In terms of the belief function Bel, (1) and (2) can be written

$$E(u^+(\delta)) = \sum_{i=1}^{n} w_i \, Bel(A_i) \qquad (3)$$

and

$$E(u^-(\delta)) = \sum_{i=1}^{n} w_i \, Bel(\overline{A_i}) \; . \qquad (4)$$

If the subsets $B_1, \ldots, B_r$ are singletons, then the belief function Bel reduces to a probability measure P. And in this case, (1) and (2) reduce, in effect, to expected utilities. Indeed if we set

$$util \, (\theta) = \Sigma\{w_i \, | \, \theta \varepsilon A_i\}$$

for all $\theta \varepsilon \Theta$, then (1) and (2) become

$$E(u^+(\delta)) = \sum_{\theta \varepsilon \Theta} P(\theta) \, util \, (\theta) = E(util(\theta))$$

and

$$E(u^-(\delta)) = \sum_{i=1}^{n} w_i - E(util(\theta)).$$

Notice that in this case the quantity $E(u^-(\delta))$ is redundant; if we are interested in comparing actions for a given situation where goals and their weights are fixed, then $E(u^-(\delta))$ only repeats the information in $E(u^+(\delta))$.

In general, if Bel is not a probability measure, then (3) and (4) depend on the full structure of the goals A_i and the weights w_i, not just on the point function $util(\delta)$. If the goals are monotonic, however, we can express $E(u^+(\delta))$ and $E(u^-(\delta))$ in terms of $util \, (\theta)$. In fact, $E(u^+(\delta))$ is the integral of $util \, (\theta)$ with respect to Bel and

$$\sum_{i-1}^{n} w_i - E(u^-(\delta)) = \sum_{i=1}^{n} w_i(1 - Bel(\overline{A_i}))$$

$$= \sum_{i=1}^{n} w_i P^*(A_i)$$

is the integral of util (θ) with respect to the <u>upper probability function</u> P^* given by $P^*(A)=1-\text{Bel}(\overline{A})$ for all A. (See Shafer 1978, pp. 3-5.)

6. Conclusion

In this paper we have seen that the theory of belief functions lends itself to a truly constructive decision theory. This decision theory is general enough to include as special cases both Bayesian expected utility calculations and the more widely used strategies of counting or weighting goals or arguments.

The new theory has several advantages over the subjective expected utility theories developed by Savage and his successors. First, it is more explicitly constructive. Since it uses the language of goals rather than the language of utility, it makes clear to the user that he must decide what to value. Secondly, it is more limited in its demand for inputs. Rather than asking a person to assign utilities to all the possibilities under consideration, it asks him only for values for the goals he cares to formulate. Rather than asking him to assign probabilities to all the possibilities, it asks him only for those probability judgments he thinks his evidence supports.

Since it is so explicitly constructive, the new theory will tend to lessen the difference that is seen between individual and social decision. Savage's picture of the ideal person, since it cannot be

applied to a group, has encouraged us to exaggerate this difference.
The new theory, since it emphasizes that basic judgments of value
and probability are themselves deliberate decisions rather than elici-
tations, can be directly used by groups as well as by individuals.

Appendix: The Problem of Small Worlds in Savage's Foundations

In this appendix I study the role of small worlds in L.J. Savage's
famous book, The Foundations of Statistics.

Savage's discussion of the problem of small worlds, on pages
82-91 of Foundations, is difficult to understand and criticize be-
cause he does not give a concrete detailed example. Here I work out
an example and use the insights gained from this example to relate
Savage's thinking to our constructive approach to decision theory.
I conclude that Savage was trapped by a picture of the "rational
person" that is more mythological than constructive, and that his
problem of small worlds dissolves when it is seen from a constructive
viewpoint. In the concluding section I argue that those of Savage's
successors who continue to call the conclusions of expected utility
theory "normative" or "prescriptive" are still trapped by Savage's
mythology.

Small Worlds

Savage formulated his postulates for expected utility theory in terms of what he called a "small world." A small world is specified by specifying two sets: a set S of alternative descriptions of the world, and a set C of descriptions of the possible consequences of certain acts. Savage called the elements of S "states of the world," and he assumed that these states or descriptions are mutually exclusive and exhaustive; the person knows that exactly one of these descriptions is true. Both S and C are relative to the person's knowledge; they are the possibilities as he sees them.

Savage assumed that S and C are formulated so that each state of the world determines a definite consequence for each of the acts the person is considering. If we denote one of the acts by f, then we can therefore think of f as a function from S to C; for each state s in S, $f(s)$ is the consequence in C that s says will result from performing the act f. Savage's postulates are concerned with the person's preferences between such acts. He showed that if the person does have well-defined preferences between such acts and these preferences satisfy certain conditions, then there exist a probability measure P on S and a real-valued function U on C such that an act f is preferred to an act g if and only if the expected value of $U(f(s))$ exceeds the expected value of $U(g(s))$.

On pages 13-15 of Foundations, Savage presented an example of a small world (S , C). He considered the situation of a person

who must decide whether to break a sixth egg into a bowl of five eggs before making an omelet. Savage supposed that the person was considering three possible acts: break the egg into the bowl, break it into a saucer, or throw it away. And he described the situation in terms of two states of the world,

$$S = \left\{ \begin{array}{l} \text{the sixth egg is good} \\ \text{the sixth egg is rotten} \end{array} \right\},$$

and six possible consequences,

$$C = \left\{ \begin{array}{l} \text{six-egg omelet} \\ \text{six-egg omelet, and a saucer to wash} \\ \text{five-egg omelet} \\ \text{five-egg omelet, and a saucer to wash} \\ \text{five-egg omelet, and one good egg destroyed} \\ \text{no omelet, and five good eggs destroyed} \end{array} \right\}.$$

Table 1, which Savage gave on page 14 of Foundations, shows how each of the three acts the person is considering can be seen as functions from S to C .

Table 1. An example illustrating acts, states, and consequences

Act	State	
	Good	Rotten
break into bowl	six-egg omelet	no omelet, and five good eggs destroyed
break into saucer	six-egg omelet, and a saucer to wash	five-egg omelet, and a saucer to wash
throw away	five-egg omelet, and one good egg destroyed	five-egg omelet

Table 1 suggests that Savage's postulates could apply to cases where relatively small numbers of acts are considered. It turns out, however, that this is not the case. The postulates require a much richer set of acts. They require, for example, that for every consequence c in C there be an act f such that $f(s) = c$ for all $s \in S$. This means that in order for the person who is considering breaking the egg to determine his probabilities and utilities, he must rank not only the three concrete acts he is considering but also many imaginary acts, such as the act that will produce a six-egg omelet no matter whether his sixth egg is good or rotten.

Savage's assumption of the availability of constant and other imaginary acts has been criticized as unrealistic. The criticism is a strong one if Savage's theory is interpreted as a descriptive theory, but it does not seem so relevant to the normative interpretation Savage preferred. We can imagine someone offering to give the person a six omelet outright, and so we can ask him (or he can ask himself) whether he would accept this offer rather than perform one of the other acts. It seemed to Savage (Foundations, page 28) that asking a person "what he would do in such and such a situation" was appropriate to his theory's "normative interpretation as a set of criteria of consistency for us to apply to our own decisions." (See also Fishburn, 1981, pp. 162-163, and Pratt, 1974, p. 88.) We can also defend such hypothetical questions using the "constructive"

vocabulary I prefer: it seems reasonable to use hypothetical choices
in the construction of preferences.

It will be useful, nonetheless, to distinguish the concrete
acts considered by the person from the wider class of acts, concrete
and imaginary, over which he is asked to express preferences. Presumably
the person first approaches the decision problem with the concrete
acts in mind. He then formalizes the problem by formulating the sets
S and C in such a way that these acts can be understood as
amounting to functions from S to C .

Refining Small Worlds

How is one small world related to a larger small world--one in
which states of the world and consequences are described in more
detail? Savage did not deal with this question directly. Instead
he discussed how a small world might fit into a grand world--a world
in which states of the world and consequences are described in ulti-
mate detail. His technical ideas about how a small world is related
to the grand world can be discussed, however, in terms of the relation
between two small worlds. By discussing them in these terms, we can,
I think, separate the essential issues from the conundrums raised
by the idea of a grand world.

Consider, then, two small worlds (S , C) and ($S\dagger$, $C\dagger$). Let
us say that ($S\dagger$, $C\dagger$) is a <u>refinement</u> of (S , C) if the following
conditions are satisfied:

(1) the elements of S correspond to a disjoint partition

of $S\uparrow$, and

(2)   the elements of C correspond to acts in the small

world ($S\uparrow$, $C\uparrow$) --i.e., each element c of C is a

function from $S\uparrow$ to $C\uparrow$.

Condition (1) means that each description in $S\uparrow$ agrees with each

description in S but goes into more detail. Condition (2) means

that a consequence in C turns out, when examined from the more

detailed perspective of ($S\uparrow$, $C\uparrow$), to mean different things depending

on which element of $S\uparrow$ is the true state of the world.

An act f in the small world (S ,C) corresponds to an act

in ($S\uparrow$, $C\uparrow$), namely, the act $f\uparrow$ whose value at an element $s\uparrow$ of

$S\uparrow$ is given by $(f(s(s\uparrow)))(s\uparrow)$. (Here $s(s\uparrow)$ denotes the element of

S corresponding to the element of the disjoint partition of $S\uparrow$

that contains $S\uparrow$.) If f formalizes a concrete act in (S,C), then

$f\uparrow$ will formalize this same concrete act in $(S\uparrow,C\uparrow)$.

We can illustrate these ideas by studying a refinement of the

small world (S , C) in Table 1. Suppose the person making the

omelet intends to serve it to a group who can distinguish between a

Nero Wolfe omelette--one made with eggs less than 36 hours old--and

an ordinary omelet--one made with eggs that are not so fresh. In

order to take this aspect of the matter into account, we refine the

states of the world to take the freshness of the eggs into account,

and we refine the consequences to take the quality of the omelet into

account. Let us suppose, for simplicity, that the person knows that

the five eggs in the bowl are all of similar freshness and that the

Table 2. The acts of Table 1, applied to a larger small world $(S\dagger, C\dagger)$

Act	State			
	Good		Rotten	
	Fresh	Stale	Fresh	Stale
break into bowl	six-egg Nero Wolfe omelet	six-egg ordinary omelet	no omelet, and five good eggs destroyed	no omelet, and five good eggs destroyed
break into saucer	six-egg Nero Wolfe omelet, & a saucer to wash	six-egg ordinary omelet, & a saucer to wash	five-egg Nero Wolfe omelet, & a saucer to wash	five-egg ordinary omelet, & a saucer to wash
throw away	five-egg Nero Wolfe omelet, & 1 good egg destroyed	five-egg ordinary omelet, & 1 good egg destroyed	five-egg Nero Wolfe omelet	five-egg ordinary omelet

sixth egg, if it is good, will not affect whether the omelet meets
Nero Wolfe standards. In this case we can describe the decision
problem in terms of a set $S\dagger$ consisting of four states of the world,

$$S\dagger = \left\{ \begin{array}{l} \text{the sixth egg is good, the other five are fresh} \\ \text{the sixth egg is good, the other five are stale} \\ \text{the sixth egg is rotten, the other five are fresh} \\ \text{the sixth egg is rotten, the other five are stale} \end{array} \right\},$$

and a set $C\dagger$ consisting of the eleven consequences listed in Table 2.

It is obvious how $S\dagger$ corresponds to a disjoint partition of S.
The two elements of S correspond to the two subsets

$$\left\{ \begin{array}{l} \text{the sixth egg is good, the other five are fresh} \\ \text{the sixth egg is good, the other five are stale} \end{array} \right\}$$

and

$$\left\{ \begin{array}{l} \text{the sixth egg is rotten, the other five are fresh} \\ \text{the sixth egg is rotten, the other five are stale} \end{array} \right\}$$

of $S\uparrow$.

The relation of C to ($C\uparrow$, $S\uparrow$) is more subtle. In order to see the consequences in C as functions from $S\uparrow$ to $C\uparrow$, we must recognize, or agree, that the omelets mentioned in these consequences will be Nero Wolfe omelets if the true state of the world in $S\uparrow$ specifies that the five eggs are fresh and an ordinary omelet other-wise. This means, for example, that the first consequence in C , c_1 = {six-egg omelet} , maps $S\uparrow$ to $C\uparrow$ as follows:

c_1(good, fresh) = six-egg Nero Wolfe omelet

c_1(good, stale) = six-egg ordinary omelet

c_1(rotten, fresh) = six-egg Nero Wolfe omelet

c_1(rotten, stale) = six-egg ordinary omelet.

One of the consequences in C appears unrefined in $C\uparrow$; this is the consequence "no omelet, and five good eggs destroyed." This consequence, conceived of as a function from $S\uparrow$ to $C\uparrow$ is, of course, a constant function; it maps all four elements of $S\uparrow$ to the consequence "no omelet, and five good eggs destroyed" in $C\uparrow$.

The equation

$$c_1(\text{rotten, fresh}) = \text{six-egg Nero Wolfe omelet} \qquad (5)$$

brings out an aspect of the relation between a small world (C , S)
and a refinement ($C\uparrow$, $S\uparrow$) that is not obvious when these small
worlds are described in terms of concrete acts, as they are in
Table 1 and 2. These tables make it clear that the "six-egg omelet"
resulting from the act "break into bowl" when the sixth egg is good
is either a Nero Wolfe omelet or an ordinary omelet, depending on
whether the other five eggs are fresh. But they do not make it clear
that the freshness of the other five eggs has any relevance to the
meaning of "six-egg omelet" in the case where the sixth egg is rotten.
Indeed, we might be inclined to doubt that it should have any relevance
in this case. If the sixth egg is rotten, then the consequence "six-
egg omelet" is realized only by an imaginary act, one that goes out-
side the factual constraints of the situation. And if we are instructed
to imagine getting a six-egg omelet in spite of the rottenness of
the sixth egg, we might think that we are also to ignore, in our imagin-
ation, the quality of the other five eggs. Equation (5) brings out that
this is not the case. Under a state of the world where a given conse-
quence is available only through an imaginary act, the meaning of the
consequence is supposed to be affected by not yet specified details
in the same way that it would be affected under a state of the world
where the consequence is possible. (More technically, the point is
that a consequence in c must be understood as mapping each state $s\uparrow$ in

$S\dagger$ to an element of C , even if s$\dagger$ cannot be associated with c by a concrete act; otherwise we cannot explain in terms of ($C\dagger$, $S\dagger$) what is meant by an imaginary act that maps s(s$\dagger$) to c.) Perhaps the possibility of missing this point when thinking about imaginary acts justifies some of the criticisms of these acts that we mentioned in the preceding section.

In the example just studied, both S and C are refined. We can construct other examples where only one of these sets is refined. If we refine S by introducing details that do not affect the consequences of the concrete acts we are considering, then it will not be necessary to refine C . And if we refine C by introducing details that are already determined for the concrete acts by the states in S , then it will not be necessary to refine C .

Refining S without refining C is of some importance in Savage's discussion. The small world (S , C) given by Table 1 does not, as it turns out, provide an adequate setting for Savage's postulates, for it is not large enough to permit as many acts as these postulates require. It can be made large enough by refinning only S . We may, for example, use a refinement $S\dagger$ of S in which every element s$\dagger$ specifies the outcome of a particular sequence of tosses of a fair coin. The outcome of such tosses will not affect the consequences of the concrete acts in the table, but the imaginary acts mapping $S\dagger$ to C will be rich enough a class for Savage's postulates.

Savage's Problem of Small Worlds

Consider a small world (S , C) and a refinement ($S\uparrow$, $C\uparrow$). Suppose a person has preferences over acts in ($S\uparrow$, $C\uparrow$) that satisfy Savage's postulates and hence determine a probability measure $P\uparrow$ on $S\uparrow$ and a utility function $U\uparrow$ on $C\uparrow$. From these preferences, probabilities, and utilities, how do we find the person's probability measure P and utility function U for the smaller world (S , C)? There seem to be two methods.

(1) Since S amounts to a disjoint partition of $S\uparrow$, we can take P to be $P\uparrow$'s marginal on that disjoint parti- tion. In symbols, $P(s) = P\uparrow(\{s\uparrow | s(s\uparrow) = s\})$. And we can say that the person's utility for a consequence c in C is his expected utility for that consequence, regarded as an act in ($S\uparrow$, $C\uparrow$). In symbols, $U(c) = E\uparrow(U\uparrow(c(s\uparrow)))$.

(2) Since every act in (S, C) can be identified with an act in ($S\uparrow$, $C\uparrow$), the person's preferences over acts in ($S\uparrow$, $C\uparrow$) determine preferences over acts in (S, C). If these latter preferences satisfy Savage's postulates, then they directly determine a probability measure P on S and a utility function U on C .

For Savage, the problem of small worlds was that these two methods may fail to produce the same answer. They will, it turns out, give the same utility function on C . But they may give different probability measures on S . (Foundations, pages 88-90.)

We can illustrate Savage's problem using the example of the omelet. Suppose a person's preferences over the small world ($S\dagger$, $C\dagger$) of Table 2 satisfy Savage's postulates and yield the probabilities and utilities shown in Table 3. (In order for Savage's postulates to be fully satisfied and the probabilities and utilities to be fully determined, we would need to further refine $S\dagger$ so that each state specifies the outcome, say, of a sequence of coin tosses. But we need not make such a refinement explicit here.) According to the probabilities in this table, the sixth egg is as likely to be rotten as good, but its being good makes it more likely that the other five are fresh. The utilities indicate that the person is indifferent to whether or not he washes a saucer or destroys a good egg, but that he prefers a six-egg omelet to a five-egg one and a Nero Wolfe omelet to an ordinary one.

Table 3. P and U for the small world ($S\dagger$, $C\dagger$) of Table 2.

states	probabilities	consequences	utilities
good, fresh	$\frac{3}{8}$	no omelet	0
good, stale	$\frac{1}{8}$	five-egg ordinary omelet	8
rotten, fresh	$\frac{1}{4}$	five-egg Nero Wolfe omelet	16
rotten, stale	$\frac{1}{4}$	six-egg ordinary omelet	16
		six-egg Nero Wolfe omelet	32

(The consequences were described in more detail in Table 2, but
we assume, for simplicity, that this further detail does not affect
the utilities. The person assigns utility 2, for example, to both
"five egg ordinary omelet" and "five-egg ordinary omelet, and a
saucer to wash.")

As it turns out the preferences over acts in the smaller world
(S, C) determined by Table 3's probabilities and utilities for
($S\uparrow$, $C\uparrow$) do satisfy Savage's postulates, and so we can apply both
methods (1) and (2) to obtain probabilities and utilities for ($S\uparrow$,$C\uparrow$).
The results are shown in Table 4. Only one set of utilities is given
in Table 4; as we have already mentioned, Savage showed that when
method (2) works it necessarily gives the same utilities as method
(1). But method (2) gives a different probability for the sixth egg's
being good than method (1) does.

The reader can easily check the numbers given in the table for
method (1). We add the probabilities $\frac{3}{8}$ and $\frac{1}{8}$ from Table 3 to obtain
$\frac{1}{2}$ as the probability of the sixth egg's being good. And the utility
calculations are straight forward. For example,

U(five-egg omelet) = $E\uparrow(U\uparrow($"five-egg omelet" $(s\uparrow)))$

$\qquad\qquad\qquad$ = $P\uparrow$(fresh)$P\uparrow$ (five-egg Nero Wolfe omelet)

$\qquad\qquad\qquad$ + $P\uparrow$(stale)$P\uparrow$ (five-egg ordinary omelet)

$\qquad\qquad\qquad$ = $\frac{5}{8}$. 16 + $\frac{3}{8}$. 8 = 13.

The reader can check that method (2) gives $\frac{7}{13}$ as the probability for
the sixth egg's being good by applying formula (7) on page 88 of

Foundations.

Why do the person's perferences over (S ,C) suggest that he assigns probability greater than $\frac{1}{2}$ to the sixth egg's being good? Because an omelet is valued more highly when the eggs are fresh than when they are stale. The distinction between fresh and stale cannot be expressed in (S ,C), but since the five eggs are more likely to be fresh when the sixth is good, the preference for fresh over stale shows up as a preference for an act that gives an omelet when the sixth is good over an act that gives an omelet when the sixth is rotten. This gives the impression that the person puts a higher probability on its being good.

Table 4. P and U for the small world (S , C) of Table 1.

states	probabilities	
	method (1)	method (2)
good	$\frac{1}{2}$	$\frac{7}{13}$
rotten	$\frac{1}{2}$	$\frac{6}{13}$

consequences	utilities (both methods)
no omelet	0
five-egg omelet	13
six-egg omelet	26

(Again, we list only three consequences instead of the six given in Table 1 because the additional detail happens to be irrelevant. Both "five-egg omelet" and "five-egg omelet, and a saucer to wash" are assigned utility 13.)

To see why P(good) appears to be equal exactly to $\frac{7}{13}$,

consider the acts f and g in (S , C) given by

    f(good) = five-egg omelet, f(rotten) = no omelet,

    g(good) = no omelet,   g(rotten) = five-egg omelet.

In terms of ($S\uparrow$, $C\uparrow$), these two acts become f↑ and g↑ ,

respectively, where

f↑(good,fresh) = five-egg Nero Wolfe omelet, f↑(rotten,fresh) = no omelet,

f↑(good,stale) = five-egg ordinary omelet,   f↑(rotten,stale) = no omelet,

g↑(good,fresh) = no omelet, g↑(rotten,fresh) = five-egg Nero Wolfe omelet,

g↑(good,stale) = no omelet, g↑(rotten,stale) = five-egg ordinary omelet.

From table 3, we obtain

E↑(U↑(f↑(s↑))) = P↑(good,fresh)U↑(five-egg Nero Wolfe omelet)

            + P↑(good,stale)U↑(five-egg ordinary omelet)

            + P↑(rotten,fresh)U↑(no omelet)

            + P↑(rotten,stale)U↑(no omelet)

$$= \frac{3}{8} \cdot 16 + \frac{1}{8} \cdot 8 + \frac{1}{4} \cdot 0 + \frac{1}{4} \cdot 0 = 7,$$

and

$$E\uparrow(U\uparrow(g\uparrow(s\uparrow))) = \frac{3}{8} \cdot 0 + \frac{1}{8} \cdot 0 + \frac{1}{4} \cdot 16 + \frac{1}{4} \cdot 8 = 6 .$$

This shows that the person prefers f↑ to g↑ , or, in the language

of the smaller world (S , C), f to g . This preference will

not be captured if we set P(good) equal to $\frac{1}{2}$ for then

$$E(u(f(s)))\ = E(U(g(s))) = \frac{1}{2} \cdot 0 + \frac{1}{2} \cdot 13 = 6\frac{1}{2} .$$

But if we set P(good) equal to $\frac{7}{13}$, we obtain expected utilities that

agree with the expected utilities for f↑ and g↑ :

$$E(U(f(s))) = \frac{7}{13} \cdot 13 + \frac{6}{13} \cdot 0 = 7$$

and

$$E(U(f(s))) = \frac{7}{13} \cdot 0 + \frac{6}{13} \cdot 13 = 6 .$$

As it turns out, this also works for all other acts in (S , C);
if we set P(good) equal to $\frac{7}{13}$, the expected utilities we calculate
for these acts in (S , C) will be the same as their expected
utilities calculated in ($S↑$, $C↑$).

The reader may be surprised that the probabilities in (S , C)
can be adjusted so as to make all the expected utilities calculated
in (S , C) come out correctly, and in fact this is possible only
because of the particular structure of the utilities in Table 2--the
Nero Wolfe omelet is always valued twice as highly as the ordinary
omelet. If we change the utility for the six-egg Nero Wolfe omelet
from 32 to 31, say, then it will no longer be possible to define
probabilities and utilities in (S , C) so as to make all the expected
utilities calculated there correct. Instead, we will find that the
person's preferences over acts in (S , C) do not satisfy the
postulates.

Savage was not disturbed by the possibility that preferences over
acts in a small world may fail to satisfy his postulates, for this
can be taken as a signal that the small world needs to be refined.
But he was disturbed by examples such as the one we have just presented.
They suggested to him the possibility that we may get a person's proba-
bilities wrong because we do not realize that we need to refine the
small world we are considering. If we check all the postulates and
find they are satisfied, how can we know that further refinement is

needed in order to make the probabilities we deduce from the person's preferences correct? Savage thought that he would probably be able to tell whether refinement is needed, but he was troubled that he could not say how to tell.

He expressed this as follows: "I feel, if I may be allowed to say so, that the possibility of being taken in by a pseudo-microcosm that is not a real microcosm is remote, but the difficulty I find in defining an operationally applicable criterion is, to say the least, ground for caution." (p. 90) By "pseudo-microcosm," he meant a small world for which a person's preferences satisfy his postulates. By "real microcosm," he meant a pseudo-microcosm for which further refinement will not show the probabilities to be incorrect.

Savage did not give a concrete example of a pseudo-microcosm that is not a real microcosm. (He did give a formal, mathematical example; see p. 89 of Foundations.) But in the light of the example we have just studied, his feeling that he would not be taken in seems like under-statement, and his desire for an "operationally applicable criterion" seems to verge on silliness. How could the person with the preferences we have been studying be unaware that these preferences depend on consideration of whether the five eggs are fresh or stale? How, in fact, could he correctly work out his attitude towards an imaginary act that gives a six-egg omelet if the sixth egg is rotten without reminding himself that ordinary, as opposed to Nero Wolfe, quality is supposed to be most likely in this case? What is the point of

asking for an operationally applicable criterion to tell him (or us) that he is taking this into account?

Myth vs. Blueprint

Savage paints a picture in which we ask a person about his preferences between certain acts. We check these preferences against a set of postulates that embody a certain notion of rationality. The preferences satisfy the postulates, at least after some minor adjustments that the person is willing to make because the postulates are so appealing. They therefore uniquely determine probabilities and utilities for the small world over which the acts are defined. These probabilities and utilities, since they are defined in terms of preferences, have an operational, behavioristic meaning. There is only one small problem: there is a remote possibility that refining the small world may show the probabilities to be erroneous.

I think it is fair to say that this picture is mythological rather than constructive. Savage's work is of great historical significance because he created a myth that could compete, within the positivist tradition, with the frequentist's myth of objective probability. Savage's myth, the myth of the rational person whose preferences can be determined operationally and used to define his probabilities and utilities, succeeded in breaking objective probability's monopoly on the imagination of the statistical community. But the myth does not work as a blueprint for constructing probabilities and utilities.

In §1 above I argued that Savage's viewpoint does not give us

a method for constructing probabilities because it goes in the wrong direction, from preferences to probabilities and utilities instead of the other way. Preferences are complicated structures--the result of construction, not the starting point.

Having studied in detail a concrete example of a small world, we can now give a deeper explanation of why preferences cannot serve as a starting point. They cannot serve as a starting point because they cannot be formulated until after the small world (S, C) is formulated, and this formulation already involves judgments of utility or, in the language I prefer, decisions on values.

Consider a person who begins with a set A_o of concrete potential acts among which he wants to choose. How does he formulate an appropriate small world? This question cannot be fully answered, for the formulation of a small world is a creative act. It is clear, however, that a person can successfully construct S and C only if he bears in mind two conditions he is trying to satisfy:

(1) Each act f in A_o and each state s in S together determine a unique element, say $f(s)$, of C, which is the consequence, as the person sees it, of the act f if s is the true state of the world.

(2) The consequences capture all the values the person brings to bear on the decision problem. This means that the person does not value any of the acts or states for their own sake. He values them only because of the consequences they deter-

mine. It also means that the level of detail of the description in C is adequate to distinguish all the values the person wants to bring to bear on the acts in A_o.

Savage states condition (1) very clearly on page 15 of <u>Foundations</u>. He never states condition (2) quite so clearly.

Aspects of condition (2) emerge when Savage presents his postulates. Consider, for example, his third postulate. In its simplest form, this postulate concerns four acts, say f,g,f', and g'. We single out consequences c_1, c_2, and c_3 and a state of the world s_o. We let f and g be constant acts,

$$f(s) = c_1 \quad \text{and} \quad g(s) = c_2 \text{ for all } s \in S ,$$

and we set

$$f'(s) = \begin{cases} c_1 & \text{if } s = s_o \\ c_3 & \text{otherwise} \end{cases} \quad \text{and} \quad g'(s) = \begin{cases} c_2 & \text{if } s = s_o \\ c_3 & \text{otherwise} \end{cases} .$$

The principle says that if g is preferred to f , then g' is preferred to f' . The underlying thought is that a preference for c_2 over c_1 should not be reversed by the specification of s_o as the state of the world. Savage explains this postulate in therms of a decision to buy a tennis racket or a bathing suit in preparation for a picnic. He points out that the postulate might be violated if we took c_1, c_2 and c_3 to be possession of a tennis racket, or a bathing suit, or nothing, respectively, and took s_o to say that the picnic would be held near a tennis court far from water. Whereas the person

might bet on the bathing suit if the location of the picnic is
unspecified, he will opt for the tennis racket if it is to be
supplied only in case there is no swimming.

> But under the interpretation of "act" and "consequence"
> I am trying to formulate, this is not the correct analysis
> of the situation. The possession of the tennis racket and
> the possession of the bathing suit are to be regarded as acts,
> not consequences. (It would be equivalent and more in accord-
> ance with ordinary discourse to say that the coming into
> possession, or the buying, of them are acts.) The conse-
> quences relevant to the decision are such as these: a refresh-
> ing swim with friends, sitting on a shadeless beach twiddling
> a brand-new tennis racket while one's friends swim, etc.
>
> (Foundations, page 26.)

It seemed clear to Savage that we could always describe "circumstances"
in such a way as to disentangle our preferences among these conse-
quences from the kind of dependence on the states involved in a
violation of his third postulate.

But how do we find the right description of consequences? How
do we find a small world (S , C) while the third postulate is satis-
fied? Do we cast about at random for ways to describe consequences
until we find one where our preferences happen to satisfy this and
the other postulates? Surely not. Surely we need to be guided by
an idea of utility or value. We need to decide what to value in the

matter at hand, so that we can deliberately try to satisfy condition
(2).

When we do succeed in satisfying condition (2), we have done more
than satisfy Savage's postulates. We have also made sure that our
small world is a real microcosm--i.e., that our values for consequences
are not dependent on the states as they were in the small world
(S , C) for the omelet. This is why Savage's call for an operationally
applicable cirterion for real microcosms is silly. We cannot satisfy
his postulates by treating them as operationally applicable cirteria.
We can satisfy them only by working in the non-behavioristic language
of condition (2). And once we permit ourselves to use this language
we see that the separation of values from states needed to satisfy the
postulates also assures that we have a real microcosm.

One aspect of Savage's discussion in Foundations that helps
sustain his behaviorism is his consistent distinction between the
person with preferences and the person who is studying them. To
use a constructive decision theory, we think about our own opinions.
But Savage's theory seems to be about how to ferret out someone else's
opinions. The two projects are not, of course, inconsistent. Perhaps
a "subject" deliberately constructs his small world, values, and
probabilities; then an "investigator" uses Savage's operationally
applicable criteria to find them out. But this picture looks frivolous
when we realize that the investigator needs to know exactly what small
world the subject constructed. From the investigator's point of view

the "problem of small worlds" is not just that an apparently
appropriate small world can turn out to be a pseudo-microcosm.
It is the more fundamental fact that he cannot get started unless he
asks the subject to tell him the small world. And if he asks for this,
he may as well ask for the other results of the construction.

The Grand World

Savage couched his discussion of the problem of small worlds
in terms of the relation between a small world and a putative
"grand world," one which has complete descriptions of the world as its
states and ultimate goods as its consequences. I have focused instead
on the relation between two small worlds (S , C) and ($S\uparrow$, $C\uparrow$)
because I think this context lends itself better to serious discussion.
The idea of a grand world leads us quickly into conundrums, and
these can distract attention from the serious issues.

One way to criticize the idea of a grand world is to say that
we cannot think of the states of a grand world as value-neutral.
These states, since they are supposed to describe the world completely
enought to determine the consequences of all acts, will have to say
whether the person will experience a long life, an enduring marriage,
or a catastrophic war. (These aspects of the state of the world may
be needed, for example, to determine the consequences of buying life
insurance.) It would even have to describe the life, marriage, or
war in detail. And we can treat such detailed descriptions as value-

neutral only by formulating very abstract "consequences" to bear the values we would ordinarily attach to longevity, fidelity, and peace. We have to talk about the "blessings of a long life," the "satisfaction of enduring marriage," and the "value of peace," and we have to think of these blessings as spiritual goods that might be bestowed on a person even if he deserts his spouse and dies young in a thermonuclear war that destroys the earth. Few moderns can countenance so spiritual a hedonism. (Examples such as longevity, fidelity, and peace suggest that it may not always make sense to separate consequences from states of the world even for small worlds. This is one aspect of the case against Savage made by Richard Jeffrey in his Logic of Decision (1965). See also Balch and Fishburn, 1974, and Jeffrey, 1974.)

Our constructive viewpoint allows us, of course, to reject the idea of a grand world without struggling with conundrums. We can simply say that all worlds are small because they have to be constructed. The detail of our descriptions of the world is limited because our time and imagination is limited.

Savage himself was uncomfortable with the idea of a grand world. He called it "unrealistic" and tried to excuse his use of it as "tongue-in-cheek" (p. 83). And he did not conceal his embarrassment that his condition for a pseudo-microcosm to be a real microcosm took the grand world "much too seriously" (p. 90). But he was forced to use the idea because he did not know how else to define a real microcosm. He did

not know how else to formulate the idea that the probabilities deduced
from preferences over acts in a small world are correct. If two small
worlds give different probabilities, which is correct? We cannot
say that the more refined is necessarily correct, for yet further
refinement might produce yet different probabilities.

Our constructive view avoids this puzzle, too. The correct
small world is the one we used to construct our probabilities and values.
There is no appeal to higher authority. Someone may convince us
to refine our small world to take into account another argument or
another item of evidence, and this may lead to different probabilities.
But this means we are changing our probabilities, not that we had
gotten them wrong.

Are Savage and His Successors Really Constructivists?

In the preceding pages I repeatedly contrasted Savage's approach
with the constructive approach that attempts first to construct prob-
abilities and utilities and then uses these probabilities and utilities
to construct preferences between acts. Savage's theory, I have suggested,
goes in the opposite direction: it shows how probabilities and utilities
can be obtained from preferences between acts.

Some readers may object to this contrast. They may argue that
Savage's postulates were only meant to describe the properties that
we would like our preferences to have, and that the practical conclusion
he drew from his demonstration that these properties implied the existence
of probabilities and utilities was precisely that we should construct

probabilities and utilities and then construct preferences from them.

Savage's writings provide little support for this commonsensial interpretation. He may have been willing to interpret his theory constructively, but he seems to have seen the elicitation of prefer-ences as the starting point in any construction. His "normative" message was that some of one's preferences may need to be changed in order to make all of one's preferences consistent with each other. And while he was willing to grant that a person's preferences are occasionally ill-defined, he regarded this as a marginal problem, not as an indication that preferences in general need to be constructed rather than elicited. (See p. 788 of Savage's 1971 paper on elici-tation.)

Savage's successors have done better. In the three decades since Savage wrote his Foundations, several authors have presented expected utility theories consciously designed to accomodate a constructive approach. The best of these theories is probably that of Pratt, Raiffa, and Schlaifer, (See Pratt, Raiffa, and Schlaifer, 1964, 1965, and Raiffa, 1968. Fishburn, 1981 describes this expected utility theory as the one most suited to "assessment" Pratt, 1974, discusses the goal of constructivity for expected utility theories.) In this theory, the calibration of probabilities by comparisons to chance models is made explicit, and it is suggested that utilities need to be assessed using simple hypothetical gambles rather than complex acts.

I would like to suggest, however, that even the most constructive of the expected utility theories have inherited some of the blinders of Savage's mythology. They have abandoned the pretence that a per-

son begins with preferences over complex acts, but they have retained the blind belief that chance models are always appropriate standards for comparison and calibration and that a person's basic preferences must determine a simple ordering.

Consider Raiffa's description of his constructive approach in Decision Analysis (1968, pp. 127-218):

> Nowhere in our analysis did we refer to the behavior of an "idealized, rational, economic man" who always acts in a perfectly consistent manner as if somehow there were embedded in his very soul coherent utility and probability evaluations for all eventualities. Rather, our approach has been *constructive:* We have prescribed the way in which an individual who is faced with a problem of choice under uncertainty should go about choosing an act that is consistent with his basic judgments and preferences.

Is a theory that prescribes the way to deal with a problem of choice truly constructive?

I would like to suggest that a theory of decision is not truly constructive unless it recognizes the need for judgment and choice at every level. Construction involves not only choice from a set of canonical examples but also the choice of a set of canonical examples. It involves not only decisions on particular preferences but also decisions on what structure for values is appropriate for a particular problem.

References

Balch, Michael, and Peter C. Fishburn

    1974   "Subjective expected utility for conditional primitives"

            (with discussion), pp. 57-96 of Essays on Economic

            Behavior under Uncertainty, M.S. Balch, D. L. McFadden,

            and S.Y. Wu, eds., North Holland/American Elsevier.

Bigelow, J. (Ed.)

    1887   The Complete Works of Benjamin Franklin, Vol. 4. New

            York: Putnam.

Brody, Richard A., and Benjamin I. Page

    1973   "Indifference, Alienation and Rational Decisions:

            The Effect of Candidate Evaluations on Turnout and

            The Vote," Public Choice 15, pp. 1-17.

Dawes, Robyn M.

    1979   "The robust beauty of improper linear models in

            decision making," American Psychologist 34, pp.

            571-582. Reprinted on pp. 391-407 in Judgment under

            Uncertainty: Heuristics and Biases, D. Kahneman,

            P. Slovic, and A. Tversky, eds., Cambridge University

            Press.

Dawes, Robyn M., and Bernard Corrigan

    1974   "Linear Models in Decision making," Psychological

            Bulletin 81, pp. 95-106.

Dickey, James M.

1980   "Beliefs about beliefs, a theory of stochastic

assessments of subjective probabilities," pp.

471-487 of Bayesian Statistics, Proceedings of the

First International Meeting Held in Valencia (Spain),

May 28 to June 2, 1979, J.M. Bernardo, M.H. DeGroot,

D.V. Lindley, A.F.M. Smith, eds., University Press,

Valencia, Spain. (Discussion on pp. 504-519.)

Fishburn, Peter C.

1981   "Subjective Expected Utility: A Review of Normative

Theories," Theory and Decision 13, pp. 139-199.

Hammond, K.R., and L. Adelman

1976   "Science, values, and human judgment," Science 194,

pp. 389-396.

Jeffrey, Richard C.

1965   The Logic of Decision, McGraw-Hill.

1974   "Frameworks for Preference," pp. 74-79 of

Essays on Economic Behavior under Uncertainty,

M.S. Balch, D.L. McFadden, and S.Y. Wu, eds.,

North Holland/American Elsevier.

Kadane, Joseph B., James M. Dickey, Robert L. Winkler, Wayne S.

Smith, and Stephen C. Peters

1980   "Interactive Elicitation of Opinion for a Normal Linear

Model," Journal of the American Statistical Association

75, pp. 845-854.

Kelley, Stanley, Jr.

  1983   Interpreting Elections, to appear.

Kelley, Stanley, Jr., and Thad W. Mirer

  1974   "The Simple Act of Voting," American Political

              Science Review 68, pp. 572-591.

Lindley, Dennis V., Amos Tversky, and R.V. Brown

  1979   "On the Reconciliation of Probability Assessments"

              (with discussion), Journal of the Royal Statistical

              Society, Series A 142 pp. 146-180.

Pratt, John W.

  1974   "Some comments on some axioms for decision making

              under uncertainty," pp. 82-92 of Essays on Economic

              Behavior under Uncertainty, M.S. Balch, D.L. McFadden,

              and S.Y. Wu, eds., North Holland/American Elsevier.

Pratt, John W., Howard Raiffa, and Robert Schlaifer

  1964   "The foundations of decision under uncertainty: an

              elementary exposition," Journal of the American

              Statistical Association 59, pp. 353-375

  1965   Introduction to Statistical Decision Theory, McGraw-

              Hill.

Raiffa, Howard

  1968   Decision Analysis: Introductory Lectures on Choices

              under Uncertainty, Addison-Wesley.

Savage, Leonard J.

1954 The Foundations of Statistics, Wiley. (Second

edition published in 1972 by Dover.)

1971 "Elicitation of Personal Probabilities and Expectations,"

Journal of the American Statistical Association 66, pp.

783-801. (Reprinted on pp. 565-614 of The Writings

of Leonard Jimmy Savage - A Memorial Selection,

published in 1981 by The American Statistical Association

and the Institute of Mathematical Statistics.)

Shafer, Glenn

1976 A Mathematical Theory of Evidence, Princeton.

1978 "Dempster's rule of combination," unpublished.

1981 "Constructive Probability," Synthese 48, pp. 1-60.

Shafer, Glenn, and Amos Tversky

1983 "The design and comparison of probability thought

experiments," unpublished.

Shaffer, William R.

1972 Computer Simulations of Voting Behavior Oxford

University Press: New York.

Wainer, Howard

1976 "Estimating coefficients in linear models: It don't

make no nevermind," Psychological Bulletin 83, pp.

312-317.