

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- 1978 "Dempster's rule of combination," unpublished.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- 1983 "The design and comparison of probability thought experiments," unpublished.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.

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.

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.

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 probabilities 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

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 enough 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-

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 accordance with ordinary discourse to say that the coming into possession, or the buying, of them are acts.) The consequences relevant to the decision are such as these: a refreshing 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 consequences 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 satisfied? 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

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_0 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_0 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-

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 understatement, 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

To see why $P(\text{good})$ appears to be equal exactly to $\frac{7}{13}$, consider the acts f and g in (S, C) given by

$$f(\text{good}) = \text{five-egg omelet}, \quad f(\text{rotten}) = \text{no omelet},$$

$$g(\text{good}) = \text{no omelet}, \quad g(\text{rotten}) = \text{five-egg omelet}.$$

In terms of (S^\uparrow, C^\uparrow) , these two acts become f^\uparrow and g^\uparrow , respectively, where

$$\begin{aligned} f^\uparrow(\text{good}, \text{fresh}) &= \text{five-egg Nero Wolfe omelet}, & f^\uparrow(\text{rotten}, \text{fresh}) &= \text{no omelet}, \\ f^\uparrow(\text{good}, \text{stale}) &= \text{five-egg ordinary omelet}, & f^\uparrow(\text{rotten}, \text{stale}) &= \text{no omelet}, \\ g^\uparrow(\text{good}, \text{fresh}) &= \text{no omelet}, & g^\uparrow(\text{rotten}, \text{fresh}) &= \text{five-egg Nero Wolfe omelet}, \\ g^\uparrow(\text{good}, \text{stale}) &= \text{no omelet}, & g^\uparrow(\text{rotten}, \text{stale}) &= \text{five-egg ordinary omelet}. \end{aligned}$$

From table 3, we obtain

$$\begin{aligned} E^\uparrow(U^\uparrow(f^\uparrow(s^\uparrow))) &= P^\uparrow(\text{good}, \text{fresh})U^\uparrow(\text{five-egg Nero Wolfe omelet}) \\ &\quad + P^\uparrow(\text{good}, \text{stale})U^\uparrow(\text{five-egg ordinary omelet}) \\ &\quad + P^\uparrow(\text{rotten}, \text{fresh})U^\uparrow(\text{no omelet}) \\ &\quad + P^\uparrow(\text{rotten}, \text{stale})U^\uparrow(\text{no omelet}) \\ &= \frac{3}{8} \cdot 16 + \frac{1}{8} \cdot 8 + \frac{1}{4} \cdot 0 + \frac{1}{4} \cdot 0 = 7, \end{aligned}$$

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^\uparrow to g^\uparrow , or, in the language of the smaller world (S, C) , f to g . This preference will not be captured if we set $P(\text{good})$ 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 $P(\text{good})$ equal to $\frac{7}{13}$, we obtain expected utilities that

(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†, C†) do satisfy Savage's postulates, and so we can apply both methods (1) and (2) to obtain probabilities and utilities for (S† , C†). 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,

$$\begin{aligned} U(\text{five-egg omelet}) &= E^\uparrow(U^\uparrow(\text{"five-egg omelet"} (s^\uparrow))) \\ &= P^\uparrow(\text{fresh})P^\uparrow(\text{five-egg Nero Wolfe omelet}) \\ &\quad + P^\uparrow(\text{stale})P^\uparrow(\text{five-egg ordinary omelet}) \\ &= \frac{5}{8} \cdot 16 + \frac{3}{8} \cdot 8 = 13. \end{aligned}$$

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

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 partition. 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.)

The equation

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

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 outside 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 imagination, 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 consequence 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

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

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^\uparrow consisting of four states of the world,

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

and a set C^\uparrow consisting of the eleven consequences listed in Table 2.

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

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 ultimate 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^\uparrow, C^\uparrow) . Let us say that (S^\uparrow, C^\uparrow) is a refinement of (S, C) if the following conditions are satisfied:

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

