

Utility, subjective probability, their interaction, and variance preferences¹

WARD EDWARDS

*Engineering Psychology Laboratory, Institute of Science and Technology,
The University of Michigan*

The Subjectively Expected Utility Maximization Model

In 1738 Daniel Bernoulli, puzzled about why people buy insurance even though insurance companies must take in more money as premiums than they pay out in benefits, proposed a theory of decision-making which asserts that people choose among risky courses of action in such a way as to maximize their expected utility (Bernoulli, 1738; Sommer, 1954). The word utility is a name for the concept of subjective value, which may be quite different from objective or dollar value. This theory was put into modern and sophisticated mathematical form in 1944 by von Neumann and Morgenstern in their famous

book, *Theory of Games and Economic Behavior* (von Neumann and Morgenstern, 1957). During the remainder of the '40s and the early '50s it attracted an increasing amount of theoretical attention and a small amount of experimental attention as well. Unfortunately, the theory did not hold up very well under experimental analysis. The difficulty seemed to be that people made decisions in risky situations on the basis of the probabilities as they perceived them, rather than the objective probabilities. This experimental finding, combined with the difficulties which arise for an expected utility maximization model when objective probabilities cannot be meaningfully defined, led several theorists to propose a new theory, which asserts that people choose among risky courses of action in such a way as to maximize what has come to be called subjectively expected utility. The subjectively expected utility of a course of action is defined as follows:

$$SEU = \psi_1 u_1 + \psi_2 u_2 + \dots + \psi_n u_n,$$

where there are n possible outcomes of the course of action, the first outcome has utility u_1 and subjective probability ψ_1 , and so on. Although a number of psychologists have expressed some severe reservations about the merits of this model in the face

¹ This is a slightly revised version of a paper presented at the meetings of the American Psychological Association, Chicago, September 1, 1960 as part of a symposium entitled "The effect of the value of a reinforcement on subjective probability." I am grateful to the other speakers in the symposium, especially to F. W. Irwin, L. Littig, and R. D. Luce, for comment and criticism, some of which is reflected in the revision. This work was sponsored by the United States Air Force under Contract AF49 (638)-769, monitored by the Air Force Office of Scientific Research of the Air Force Office of Aerospace Research, and carries Document No. AFOSR-1709.

of the data, it remains the dominant model for static decision-making to this day—although nowadays it is usually complicated by a stochastic formulation which causes it to predict a probability that one course of action will be preferred to another, rather than a choice. For reviews of this literature, see (Edwards, 1954d, 1961).

Non-independence of Values and Probabilities

One important feature of the subjectively expected utility maximization model is that the utilities are ordinarily assumed to be functions only of the possible outcomes of the course of action, while the subjective probabilities are ordinarily assumed to be functions only of the objective probabilities of those outcomes. Neither of these assumptions is necessary to the mathematical content of the model itself, but it is very difficult to see how the model could be applied to real decisions unless some such assumptions were made. Nevertheless, a number of people have proposed that utility is determined both by objective probability and by objective value, and that subjective probability is also determined both by objective probability and by objective value.

Specifically, two propositions of this sort have been defended. One asserts that an outcome which has a low probability will, by virtue of that low probability, have a higher value or utility than the same outcome would have if it had a high probability. The assertion that the grass is greener on the other side of the fence embodies the essence of this proposition. Since this hypothesis was originally brought to prominence in psychology by Kurt Lewin and his students, I will call it the Lewin utility theory. The other common proposition concerns the determiners of subjective probability; it says that people will overes-

timate the likeliness of desirable events and underestimate the likeliness of undesirable ones. A number of psychologists have proposed this second hypothesis in one form or another. Since Francis W. Irwin is an early and effective proponent of this view, I shall call it the Irwin subjective probability theory (Irwin, 1953). It is important to note that although both of these theories are inconsistent with the simple version of the SEU maximization model, they are nevertheless two completely independent theories; either one could be correct while the other was wrong.

Before evaluating the validity of the Lewin utility theory and the Irwin subjective probability theory, we must first decide what the simpler hypothesis is with which both of these theories disagree. This is not as easy as it sounds. Unfortunately, neither subjective probability nor utility is usually conceived of as directly observable, or even as easily inferred from observations. We usually suppose that choices among risky courses of action are determined by both utilities and subjective probabilities, and we then attempt to infer appropriate values of these quantities to account for observed sets of choices. Since we permit considerable arbitrariness in determining utility and subjective probability values, it takes some substantial set of observations, very carefully chosen, to obtain a contradiction of the SEU model. Without discussing the evidence in detail, I would like to assert that I know of no experiment in the psychological literature, including Irwin's recent one (personal communication), which cannot be explained in terms of appropriate utility and subjective probability functions. Consequently I am compelled to argue that no evidence against the SEU model and in favor of either the Lewin utility theory or the Irwin subjective probability theory exists.

I do not consider this fact to be in favor of the SEU model, however. What it means is that the SEU model, in the form in which it has usually been defended, is so vague that it is almost impossible to find a reasonable pattern of behavior which is inconsistent with it.² The SEU model can be made sufficiently specific so that it really means something only by specifying the forms of the utility and subjective probability functions which enter into it, or (what amounts to the same thing) by specifying some operations by which utilities and subjective probabilities can be measured independently of the decisions they are supposed to predict. Just such a development is now in progress. Direct psychophysical methods are being applied to the measurement both of utility and of subjective probability. So far the results of using such directly obtained utility and subjective probability functions to predict decisions among bets have not been very favorable to SEU models. Unfortunately, they are not very favorable to the Lewin or Irwin theories either. Such experiments seem to show that the utility of money is more or less linear with its dollar value (for small amounts of money, at any rate), and that the estimated probability associated with a suitable probability display is very close to linear with the objective probability being displayed. In short, the data indicate that people perceive both values and probabilities correctly. Nevertheless, they do not make decisions which maximize objective expected value, so it must be concluded that although they perceive values and

² Since this paper was written, I have reported (Edwards, 1962) a deduction from the SEU model so specific that it would be very easy to find data which violate it, which assumes nothing but continuity and monotonicity of utility and subjective probability functions. This implies that the SEU model is more vulnerable than I realized in 1960.

probabilities correctly, they use them incorrectly. In other words, on this interpretation it is the mathematical form of the SEU model, rather than the links between it and the real world of choices, which is at fault. (This argument is spelled out in detail in [Edwards, 1962]). I will discuss a specific hypothesis about how the mathematical form of the SEU model is at fault and how to correct it later in this paper.

Evidence Concerning Non-independence

So far I have been treating the Lewin utility theory and the Irwin subjective probability theory as though they were alternatives to the SEU maximization model, to be adopted only if the SEU maximization model turns out not to work. But they can equally well be considered as models in their own right, to be used independently of the SEU model. When they are thought of in that way, a substantial amount of evidence bearing on them is available.

First, consider the evidence for the Irwin subjective probability theory. The results of several experiments show that people consider an event more likely to occur if its consequences are favorable than if its consequences are unfavorable (Crandall, Solomon, and Kellaway, 1955; Irwin, 1953; Marks, 1951). Some of the data from my probability preference experiments indicate more or less the same thing (Edwards, 1953, 1954a, 1954b). In those experiments subjects were asked to make choices from pairs of bets. Both members of each pair had the same expected value, so there was never any objective reason for preferring one bet to the other. Figure 1 shows the results for positive expected value bets. The y -axis is a measure of the extent to which various bets are preferred, and the x -axis is the probability of winning of the bets. It is evident that people particularly like the 50-50 bet and particularly do not




FIG. 1. Preference as a function of probability for positive expected value bets. (Reproduced by permission from Edwards, 1953, p. 355).

like the 6/8 bet; these findings are highly reproducible. Figure 2 shows the same information for negative expected value bets. Here the outstanding finding is that people prefer relatively low probabilities of losing (and relatively high amounts of loss) and avoid relatively high probabilities of losing, and relatively small amounts of loss. These findings would be consistent with the hypothesis that people considered an event with a negative expected value to be less likely than the same event with a positive expected value, although of course they do not prove that hypothesis.

Another form of evidence comes from an experiment in which I actually obtained utility and subjective probability functions (Edwards, 1955). Again subjects made choices among bets, and these functions were inferred from their choices. The method used depended on a rather doubtful assumption, but the data obtained in the experiment provided internal consistency checks on the assumption which

worked pretty well, so I have some confidence in the data even though I do not like the assumption used to obtain them. The data relevant to the Irwin subjective probability theory are presented in Figure 3. Ten subjective probability functions, two for each of five subjects, are shown. For each subject, one curve is his subjective probability function inferred from positive expected value bets and the other is his subjective probability function inferred from negative expected value bets. It is apparent that most subjects remained pretty objective when choosing among negative expected value bets, but were over-optimistic when choosing among positive expected value bets.

A Conclusion about the Irwin Subjective Probability Theory

All of these findings strongly indicate that there is at least an interaction between the *sign* of the utility of a bet and the subjective probability associated with the




FIG. 2. Preference as a function of probability for negative expected value bets. (Reproduced by permission from Edwards, 1953, p. 355.)

event (note that the utility must be measured on a ratio scale for this statement to be meaningful, and zero utility must be where you now are). Furthermore, the direction of the effects is in general the direction predicted by the Irwin subjective probability theory. None of the evidence, however, indicates an interaction between value and subjective probability provided that the signs of the utilities involved do not change. This is fortunate for the SEU model. It is not too difficult to think of an SEU model which requires up to five different functions relating subjective to objective probability, depending on the signs of the utilities involved. One such function might apply when all utilities were positive, another when the lowest utility was zero, a third when both positive and negative utilities were present, a fourth when the highest

utility was zero, and a fifth when all utilities were negative. Some of the data on zero expected value bets suggests that there may be little or no difference between the first two or three of these cases. But I do not see how to formulate a model which would permit the dependence of subjective probability on amounts, as well as signs, of utilities. And I know of no evidence which requires anything more than interaction with sign. So, for the time being, perhaps we can afford to assume that the Irwin subjective probability theory is right, but that the dependence of subjective probability is on signs of utilities rather than on their values.

Atkinson's Form of the Lewin Utility Theory

Now consider the Lewin utility theory. I don't propose to examine it in Lewin's

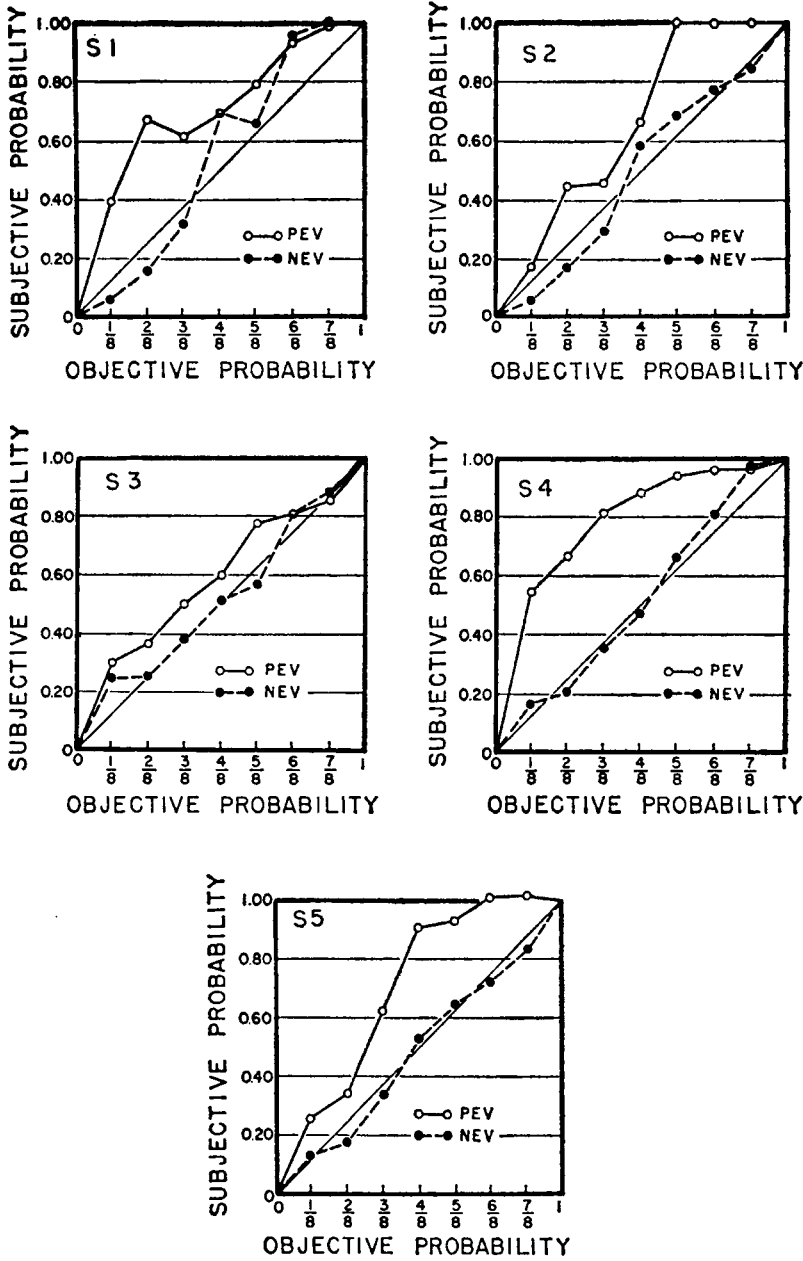


FIG. 3. Subjective probability as a function of objective probability and sign of expected value. (Reproduced by permission from Edwards, 1955, p. 208.)

form; instead, I will take the model which Atkinson has proposed as my point of departure (Atkinson, 1957; Atkinson, Bastian, Earl, and Litwin, 1960). Atkinson believes that in situations in which a person perceives the outcome of a risky act as contingent on his skill, the variables which will control his decision are the motive to achieve success (M_s), the motive to avoid failure (M_f), the subjective probabilities of achieving success and avoiding failure (P_s and P_f) and the incentive values of achieving success and avoiding failure (I_s and I_f). He combines these six determiners of decisions into a single resultant motivation by simple multiplication; his equation for resultant motivation is:

$$(M_s \cdot P_s \cdot I_s) + (M_f \cdot P_f \cdot I_f).$$

Now Atkinson's crucial assumption is that the attractiveness of success is a positive linear function of task difficulty, and similarly that the attractiveness of failure, always a minus quantity, becomes increasingly negative for easier tasks. In symbols:

$$P_f = 1 - P_s; \quad I_s = 1 - P_s; \quad I_f = -P_s$$

It follows that Atkinson's equation for resultant motivation can be rewritten as follows:

$$P_s(1 - P_s)(M_s - M_f)$$

You may already have noticed an interesting similarity between this equation for resultant motivation and a familiar statistical quantity: the variance of a two-outcome bet. If you can win \$A with probability p and win \$B with probability $1 - p$, then the variance of the bet is

$$p(1 - p)(A - B)^2.$$

The only structural difference between the equation for variance and Atkinson's equation for resultant motivation is that in Atkinson's equation the difference between

the two motives is not squared. This means that Atkinson's resultant motivation can be negative, while of course a variance cannot be less than zero.

Variance Preferences

This structural similarity suggests that Atkinson's model, though not couched in that language, is in fact a suggestion that people may in choosing among bets base their preferences not only on the first moment of the distribution of outcomes (the mean, or expected value, or SEU), but also on higher moments of the distribution. That suggestion has been made before; in fact, it has been perhaps the most persistent criticism of the school of thought about utility maximization which was started by von Neumann and Morgenstern (Allais, 1953). In 1954 I published an experiment addressed to the question, which concluded that variance preferences do indeed exist, but are minor in importance compared with utility and subjective probability as determiners of choices among bets (Edwards, 1954c). More recently, Coombs and Pruitt have performed an experiment in which they asked subjects to make choices from pairs of imaginary bets of fixed expected value which varied in variance and in skewness (Coombs and Pruitt, 1960). Data were rejected if they failed to satisfy weak stochastic transitivity, which is a fairly lenient kind of assumption that if A is preferred to B and B is preferred to C, then A ought to be preferred to C at least 50 per cent of the time. The remaining data were analyzed according to Coombs' unfolding technique. The major findings were that most subjects chose consistently and transitively and that most subjects exhibited single-peaked variance preferences and (except for an undue preference for 50-50 bets) single-peaked skewness or probability preferences. The consistency found by

Coombs and Pruitt was extraordinarily high; perhaps it is because their subjects made their choices among these imaginary bets at rates ranging from 3.4 to 22.4 choices per minute, and so cannot have taken much time for reflection. The two rules "Always choose the bet with the highest payoff for winning" and "Always choose the bet with the lowest cost for losing" together account for about 68 per cent of all the rank orderings Coombs and Pruitt obtained; these rules, of course, could be applied simply and mechanically to minimize thought. It has been shown (Slovic, Lichtenstein, and Edwards, forthcoming) that more complex preference patterns, less easily interpreted as variance preferences, occur when subjects make more careful choices, but of course this finding is not evidence against variance preferences; in fact, the evidence suggests that they do exist, but not in the single-peaked form hypothesized by Coombs and Pruitt.

Unfortunately, for any bet which has only two outcomes, skewness is completely confounded with probability of winning and variance is completely confounded with utility. Consequently any experiment designed to examine variance preferences unequivocally must use bets with more than two possible outcomes. So neither my earlier experiment nor the Coombs-Pruitt experiment nor the Slovic-Lichtenstein-Edwards experiment has anything definitive to say about the existence or relative size of variance preferences. Nevertheless, it seems intuitively plausible that such preferences exist, and Atkinson's theory can be reinterpreted very simply to be an assertion that people with high need for achievement and low need to avoid failure should prefer high-variance bets, while people with low need for achievement and high need to avoid failure should prefer low-variance

bets—a hypothesis which is supported by a substantial amount of data.

A Conclusion Concerning the Lewin Utility Theory

If variance preferences exist, then the SEU model is not very useful in predicting choices. But the notions of utility and subjective probability are necessary aside from their use in SEU models. People do make judgments about the values of things, and they do make judgments about the likelihood of events. These judgments can conveniently be called utilities and subjective probabilities, regardless of the usefulness or lack of usefulness of the SEU maximization model. But the Lewin utility theory says more than that variance preferences exist. It asserts that objects which are harder to obtain are more valuable than those which are easier to obtain. This hypothesis has meaning independent of choices among risky courses of action. But I am unclear, and have been unable to find out from the literature, which of two possible interpretations of this statement is intended by those who make it. One interpretation simply asserts that our world is so constructed that the more desirable objects are harder to get. I feel confident that this is so, but I see no reason why this fact should be interpreted as a theory of behavior. It is instead a theory about the environment in which behavior takes place. The other possible interpretation is that exactly the same object is more valuable when it is hard to obtain than when it is easy to obtain. This hypothesis is intuitively plausible, but I know of no very convincing evidence that it is correct.

Early in this paper, I said that the Lewin utility theory is summarized by the proverb "The grass is greener on the other side of the fence." I am now saying that there are two possible interpretations of this proverb.

One is that the grass is really the same on both sides, but looks greener over there. The other, which seems to me more plausible, is that the grass is in fact greener over there—and indeed they put up the fence in the first place just to keep it that way. Until evidence against the second of these two interpretations accumulates, I find it more persuasive than the first.

Summary

The subjectively expected utility maximization model, currently the dominant static model about how men make risky decisions, rests on the concepts of utility or subjective value, assumed to be a function only of amount and nature of the valued object, and subjective probability, assumed to be a function only of amount and display or means of realization of objective probability. But evidence, anecdotal and experimental, suggests that human judgments of value are affected by probability of obtaining the valued object, and human judgments of probability are affected by value of the possible payoffs of the risky course of action. Two specific theories of this sort are considered. The Irwin subjective probability theory asserts that people will overestimate the likeliness of desirable events and underestimate the likeliness of undesirable events. Review of relevant data indicates that such an interaction indeed occurs, but that it is only with the signs, not with the amounts, of the possible payoffs. The Lewin utility theory asserts that an outcome with a low probability will have a higher value than the same outcome would have if it had a high probability. A specific version of this point of view proposed by Atkinson is reviewed and found to resemble closely the hypothesis that choices among risky courses of action depend on the variance as well as the expected value of each course. Evidence for

such variance preferences is reviewed. It is concluded that the favorable evidence concerning variance preferences casts serious doubts on the validity of the subjectively expected utility maximization model. But judgments of value can occur independently of decisions, and so the Lewin utility theory is important even if the subjectively expected utility maximization model is wrong. The question is raised whether the attractiveness of improbable achievements is not a characteristic of the environment rather than of judgments of value. Doesn't the grass seem greener on the other side of the fence because it is greener, and wasn't the fence erected to keep it that way?

REFERENCES

- ALLAIS, M. "Le comportement de l'homme rationnel devant le risque: critique des postulats et axiomes de l'école américaine," *Econometrica*, 21 (1953), 269-90.
- . "La psychologie de l'homme rationnel devant le risque: la théorie et l'expérience," *Journal de la Société de Statistique de Paris*, 94 (1953), 47-73.
- ATKINSON, J. W. "Motivational Determinants of Risk-taking Behavior," *Psychological Review*, 64 (1957), 359-72.
- ATKINSON, J. W., BASTIAN, J. R., EARL, R. W., and LITWIN, G. H. "The Achievement Motive, Goal Setting, and Probability Preferences," *Journal of Abnormal and Social Psychology*, 60 (1960), 27-36.
- BERNOULLI, D. "Specimen theoriae novae de mensura sortis," *Comentarii Academiae Scientiarum Imperiales Petropolitanae*, 5 (1738), 175-92. Translation by L. SOMMER in *Econometrica*, 22 (1954), 23-36.
- COOMBS, C. H. and PRUITT, D. G. "Components of Risk in Decision-making: Probability and Variance Preferences," *Journal of Experimental Psychology*, 60 (1960), 265-77.
- CRANDALL, V. J., SOLOMON, D., and KELLAWAY, R. "Expectancy Statements and Decision Times as Functions of Objective Probabilities and Reinforcement Values," *Journal of Personality*, 24 (1955), 192-203.

- EDWARDS, W. "Probability Preferences in Gambling," *American Journal of Psychology*, 66 (1953), 349-64.
- . "Probability Preferences among Bets with Differing Expected Values," *American Journal of Psychology*, 67 (1954a), 56-7.
- . "The Reliability of Probability Preferences," *ibid.* (1954b), 68-95.
- . "Variance Preferences in Gambling," *ibid.* (1954c), 441-52.
- . "The Theory of Decision-making," *Psychological Bulletin*, 51 (1954d), 380-417.
- . "The Prediction of Decisions among Bets," *Journal of Experimental Psychology*, 51 (1955), 201-14.
- . "Behavioral Decision Theory," *Annual Review of Psychology*, 12 (1961), 473-98.
- . "Subjective Probabilities Inferred from Decisions," *Psychological Review*, 69 (1962). (In press.)
- IRWIN, F. W. "Stated Expectations as Functions of Probability and Desirability of Outcomes," *Journal of Personality*, 21 (1953), 329-35.
- MARKS, R. W. "The Effect of Probability, Desirability, and 'Privilege' on the Stated Expectations of Children," *Journal of Personality*, 19 (1951), 332-51.
- PHARES, E. J. "Expectancy Shifts in Skills and Chance Situation," *Journal of Abnormal and Social Psychology*, 54 (1957), 339-42.
- SLOVIC, P., LICHTENSTEIN, S. C., and EDWARDS, W. "Boredom: A Vitiating Variable in Decision Experiments." (Forthcoming.)
- VON NEUMANN, J., and MORGENSTERN, O. *Theory of Games and Economic Behavior*. 2nd ed. Princeton, N.J.: Princeton University Press, 1957.