JUDGEMENT AND DECISION MAKING UNDER UNCERTAINTY: WHAT ECONOMISTS CAN LEARN FROM PSYCHOLOGY

Richard Thaler

What can economists learn from psychologists in the areas of judgement and decision making? We know that economics is pretty good as a normative model - marginal cost equals marginal benefit is the right answer. But can we use standard economic theory to predict what people will do? There is a basic distinction between positive and normative theories which is intriguing and which I have been investigating from several aspects over the last four years (12, 15, 16, 17). What I would like to do in this lecture is to talk mainly about what psychologists are doing in this area and how the results of their research may be applied in the area of positive economic theory.

Consider the following question. If there are twenty-five people in a room, what is the probability that at least two will share a birthday? This is a very famous problem primarily because no one seeing it for the first time gets it right. The correct answer to the problem is a little over .5, and if you get as few as fifty people in a room the probability of at least one shared birthday approaches 1. It should be clear to both psychologists and economists that one can make quite a bit of money on bets on this proposition.

What is particularly interesting with respect to the topic I am addressing is the systematic nature of the error. If a naive subject is presented with this problem or its formal equivalent we can predict both that they will get it wrong and that the error will be in a specific direction. This is the paradigm for the research I'm going to discuss. The approach is to focus on a particular problem, identify its formal structure, analyze how people solve that problem, and observe how the reported solution differs from the correct solution. From this analysis we try to infer what they are doing and how their decision process differs from that which the normative theory would suggest.

Work in the area of behavioral decision theory can be divided into two broad categories: judgement and decision making. Judgement refers to the process of making probability estimates. For example, a statement of the probability that it will snow tomorrow is a judgement. The process of determining whether one should wear an overcoat is a decision. Two articles which survey and summarize the kinds of problems which are addressed and the work which has been done in this area are Slovic, Fischoff and Lichtenstein (13) and Tversky and Kahneman (18).

Richard Thaler is Assistant Professor in the College of Business and Public Administration, Cornell University. He was assisted by Karen Farkas in preparing the text for publication.

Tversky and Kahneman have looked at the way people make probability estimates and have argued that people use simple rules of thumb, or heuristics, to help them make judgements. They go on to point out that these heuristics have biases similar to those in the birthday problem.

Let's consider a second problem - the Tom W. problem.

> Tom W. is of high intelligence, although lacking in true creativity. He has a need for order and clarity and for neat and tidy systems in which every detail finds its appropriate place. His writing is rather mechanical, occasionally enlivened by somewhat corny puns, and by flashes of imagination of the sci-fi type. He has a strong drive for competition. He seems to have little feeling and sympathy for other people and does not enjoy interacting with others. Self-centered he nonetheless has a deep moral sense.
>
> --The above is a psychological profile of a graduate student in the U.S. He is in one of three fields: education, computer science, or the humanities. Indicate the probabilities you associate with the likelihood of his being in each of the three fields. (Hint: they should add up to 1.)

The description is clearly the stereotype computer jock. When people are asked to indicate probabilities as described, most give a very high probability to the choice of computer science. This ignores what psychologists refer to as the base rate, or what statisticians might call the prior. There are many more people in the fields of education and humanities than there are in the field of computer science. In fact, there are more people who fit the computer science stereotype but are in humanities than there are people who are in computer science and fit this stereotype.

The mistake people are making in this problem is the utilization of a heuristic which Kahneman and Tversky call representativeness; i.e., they ask themselves how representative this description is of computer scientists they have met. If the description is very similar to most (or many) computer scientists they have met, they then judge the probability that this person is a computer scientist to be high. That is in fact exactly what seems to have happened for this set of problems. Kahneman and Tversky asked one group of people how similar this description was to computer scientists they had met. They then asked a different group to determine the probability that Tom W. is a computer scientist. In both cases, they got the same answer. People seem to treat these problems as one and the same whereas they are in fact different.

They then used another set of questions, of a similar form, to study representativeness. A description of an individual would be presented to a group of people who were told that it came from a sample of one hundred individuals, of whom seventy were engineers and thirty were lawyers. The group was asked to report the probability that the description was that of one of the engineers. They then asked a different group the same question, this time reversing the proportions. Clearly, the group which is told that the sample is 70:30 engineers:lawyers should give a higher probability than should the group which is told that the sample is 70:30 lawyers:engineers. However, they gave exactly the same answer.

We can highlight this result with an additional experiment in which the following description was supplied:

> Dick is a 30-year-old man. He is married and has no children. A man of high ability and high motivation, he promises to be quite successful in his field. He is well liked by his colleagues.

Participants were then asked the probability that Dick was one of the seventy lawyers, or thirty lawyers, depending on which version of the problem had been presented. Now this is clearly a description that has no content with respect to occupation. The "correct" answer to whether he is one of the seventy (thirty) lawyers is .7(.3). The mode answer given is .5. Now, what is interesting is that if you give people no information other than the mix of lawyers and engineers, and ask the probability that a random person from the sample is an engineer, they will respond with the prior (.7 or .3). But if you give the innocuous description which should have no effect on their judgements they say .5.

This was essentially Kahneman and Tversky's first pass at dealing with Bayes' rule as a descriptive model. In Bayes' rule, equal weight is given to the prior odds and to the likelihood ratio. What seems to be going on here is that people are underweighting the prior odds and overweighting the likelihood ratios. Representativeness is one possible explanation, but another explanation which is perhaps more interesting is investigated with the following situation (19).

> A cab was involved in a hit-and-run accident at night:
> Two cab companies, the Green and the Blue, operate in the city.
> You are given the following data:
>
> i) 90% of the cabs in the city are Green and 10% are Blue.
>
> ii) A witness identified the cab as Blue. The court tested his ability to identify cabs under the appropriate visibility conditions. When presented with a sample of cabs (half of which were blue and half of which were green) the witness made correct identifications in 80% of the cases and erred in 20% of the cases.
>
> Question: What is the probability that the cab involved in the accident was Blue rather than Green?

The subjects used in this experiment were typically students. Sometimes the question was put on entrance exams for Hebrew University, which would lead one to infer that the students would be highly motivated to answer correctly.

What we have here is a classic Bayes' rule application. The prior is that 90% of the cabs are green and 10% are blue. There is also a piece of information that a witness has identified the cab as blue, and that this witness is accurate 80% of the time. If you work through the normative model you find the probability that the witness is correct is .31, which is lower than one's intuition would expect. When this question is given to groups of subjects most of them answer .8, explaining that if the witness is 80% accurate the reported observation will be correct 80% of the time.

What is going on here? The explanation that Kahneman and Tversky offer is that people think about this type of problem using what is referred to as a causal schematic. In other words, they try to put things into a deterministic framework in which there are clear cause and effect relationships. In this problem, the fact that 90% of the cabs are green does not fit into their causal schematic, and so they throw away that piece of information. There is some evidence to support the hypothesis that people can be manipulated to use this information if, instead of saying that 90% of the cabs are green, you say that 90% of the accidents involving taxi cabs are with green cabs. This change in presentation does not change the formal properties of the problem from the statistician's point of view. However, it appears that people will then regard the base rate information as having a relationship to cause and effect and will incorporate it into their probability estimation. The result is that the modal answer falls below .8, although not sufficiently below to reach the normative answer. So it appears that if you can give people a description in which the base rate seems to play a cause and effect role with respect to the likelihood ratio, the information will be used, albeit insufficiently; otherwise, it will not be used at all.

Students are not the only people who have problems in this area. For example, a farmer had some damage done to his barn as a result of something falling off a low flying plane. He sued the U.S. Air Force for damages, and testimony was introduced that 95% of the planes that flew over this farmer's land belonged to the Air Force. The judge threw out that testimony as irrelevant. Then the farmer took the stand and testified that as the plane flew over he was able to see into the cockpit and observed that the pilot was wearing an Air Force uniform. That information was accepted.

Another example occurred while I was watching a football game one afternoon. Tom Brookshire, the announcer, was talking about Tony Dorsett, and observed that for games in which Dorsett carries the ball more than twenty times Dallas' record is 15-1. This was presented as evidence that Tony Dorsett is a good football player. You might be interested in knowing that when I played on my highschool basketball team, the games in which I played were won by an average of thirty points. I'm still waiting to be honored in the Hall of Fame.

An econometrician at Cal Tech, David Grether, has replicated some of Kahneman and Tversky's work (5). What he did was to run an experiment in which he had two bingo cages, from one of which he drew data and presented it to a group of students. The problem was to determine from which cage the data had been drawn. This was another Bayes' rule application, and the same basic results were observed. Students tended to overweight the likelihood ratio and underweight the priors. In odds formulation, these should be weighted equally. His subjects gave a one-third weight to the prior and a two-thirds weight to the likelihood ratio.

Kahneman and Tversky talk about another heuristic which they call the availability heuristic. This relates to a situation in which a person has to make a judgement about how likely something is, and does so by trying to recall instances of that event. According to their theory, the easier it is to recall instances the more likely the event will be judged to be. That seems to be a pretty good heuristic. Actually, none of the heuristics discussed would have survived if they were not pretty good. What makes them interesting is that they have systematic biases. For example, consider words of three or more letters in the English language. Which do you think is more likely, words that start with the letter "r" or words that have an "r" in the third position? Well, most people think it's the former - words that start with the letter "r" - whereas in fact it is the latter. The proposed explanation for this is that it is much easier to think of words that begin with a letter than words that have that letter in the middle because your mind, faced with that kind of task, tends to work like a dictionary. It is hard to come up with a systematic way of thinking of words that have "r" in the third position. The relative ease of one task over another creates a bias in a predictable direction.

Similarly, if you ask people whether there are more homicides or suicides, most people respond that there are more homicides. In fact, there are more suicides. But it is easier to think of homicides because they get more publicity. Lichtenstein, Slovic, Fischhoff _et_ _al_. (8) have done a big study of people's perceptions of both risk and causes of death and have found substantial evidence of this phenomena. The moral of the story up to this point is that in making these kinds of judgements people use simple heuristics that are pretty good on average, but which have predictable biases leading people to make predictable mistakes. Therefore, if we want to predict behavior we should use the applicable heuristic rather than the normative model. If we wanted to predict what people would say when given the birthday problem, for example, we would not give the normative answer, but something much lower.

The next topic I want to cover is what is referred to in the literature as calibration. Calibration means that if I ask you what the probability is that the Steelers are going to win the Super Bowl, and you say .9, then we should observe that 90% of the time they win and 10% of the time they lose. If they win all the time you would have been underconfident; if they win only half the time you would have been overconfident. Putting estimated probability on the horizontal axis and observed frequency on the vertical axis, the estimates of a perfectly calibrated person ought to lie on a forty-five degree line from the origin.

There has been considerable study of how well calibrated people actually are (9). The results are that weather forecasters are very good and everyone else is pretty lousy. Referring to Figure 1, the solid forty-five degree line represents perfect calibration. The results of various studies of forecasters' responses have been indicated by different configurations, and you can see that they are all reasonably close to a forty-five degree angle. There are some deviations, but by and large they are very good. It should be noted that weather forecasters were not always this well calibrated. Probability estimates began being used to a significant extent some time in the late fifties. Initially forecasters were not very good, but as they gained esperience in making these judgements they became much better calibrated.

The results of a series of calibration studies are presented in Figure 2. Subjects were asked questions of the form, "Which city has a larger population, Portland or Tucson, and what is the probability that you are correct?" Clearly you should never say less than .5 in response to a question like this, and just to make sure no one does, they are not allowed to (otherwise they might). Subjects were given one or two hundred of this sort of question, another example of which was, "What is absinthe, a precious stone or a liquor?" You can see that although the plotted results are from four different studies, the pattern is remarkably similar. What is shown is a great degree of overconfidence on the part of the subjects. For example, look above .7 on the horizontal axis. These are the percent correct for those questions for which subjects said they would be correct 70% of the time. You can see that they actually were correct only 55-58% of the time. What is even more striking are those responses for which subjects said they were 100% certain; the proportion correct was only 80%. I have run a similar experiment with my students, and also found an 80% rate on those questions for which the students said they were sure they were right.

I'm not really sure what all of these results tell us, the main reason being that what economists typically study is decision making, not judgement. What the calibration results show is that when people are asked to give a probability, the one they give is likely to be wrong. The missing link is how these errors in judgement are incorporated in the ultimate decision. For example, there could be cancelling errors. If a person assesses a probability to be 80%, when it should in fact be 90%, but then goes on to overweight the probability in the process of decision making, it could cancel out. I think that what we basically have here is the observation that if people are asked to predict an event they will probably be overconfident - overconfident in their ability to predict - that is how I would summarize these results.

I would like to turn my attention at this time to the subject of decision making. The normative model of decision making under uncertainty is the Von Neumann-Morgenstern-Savage expected utility theory (11, 20). As many people know, that model has been questioned as to whether it is

Figure 1

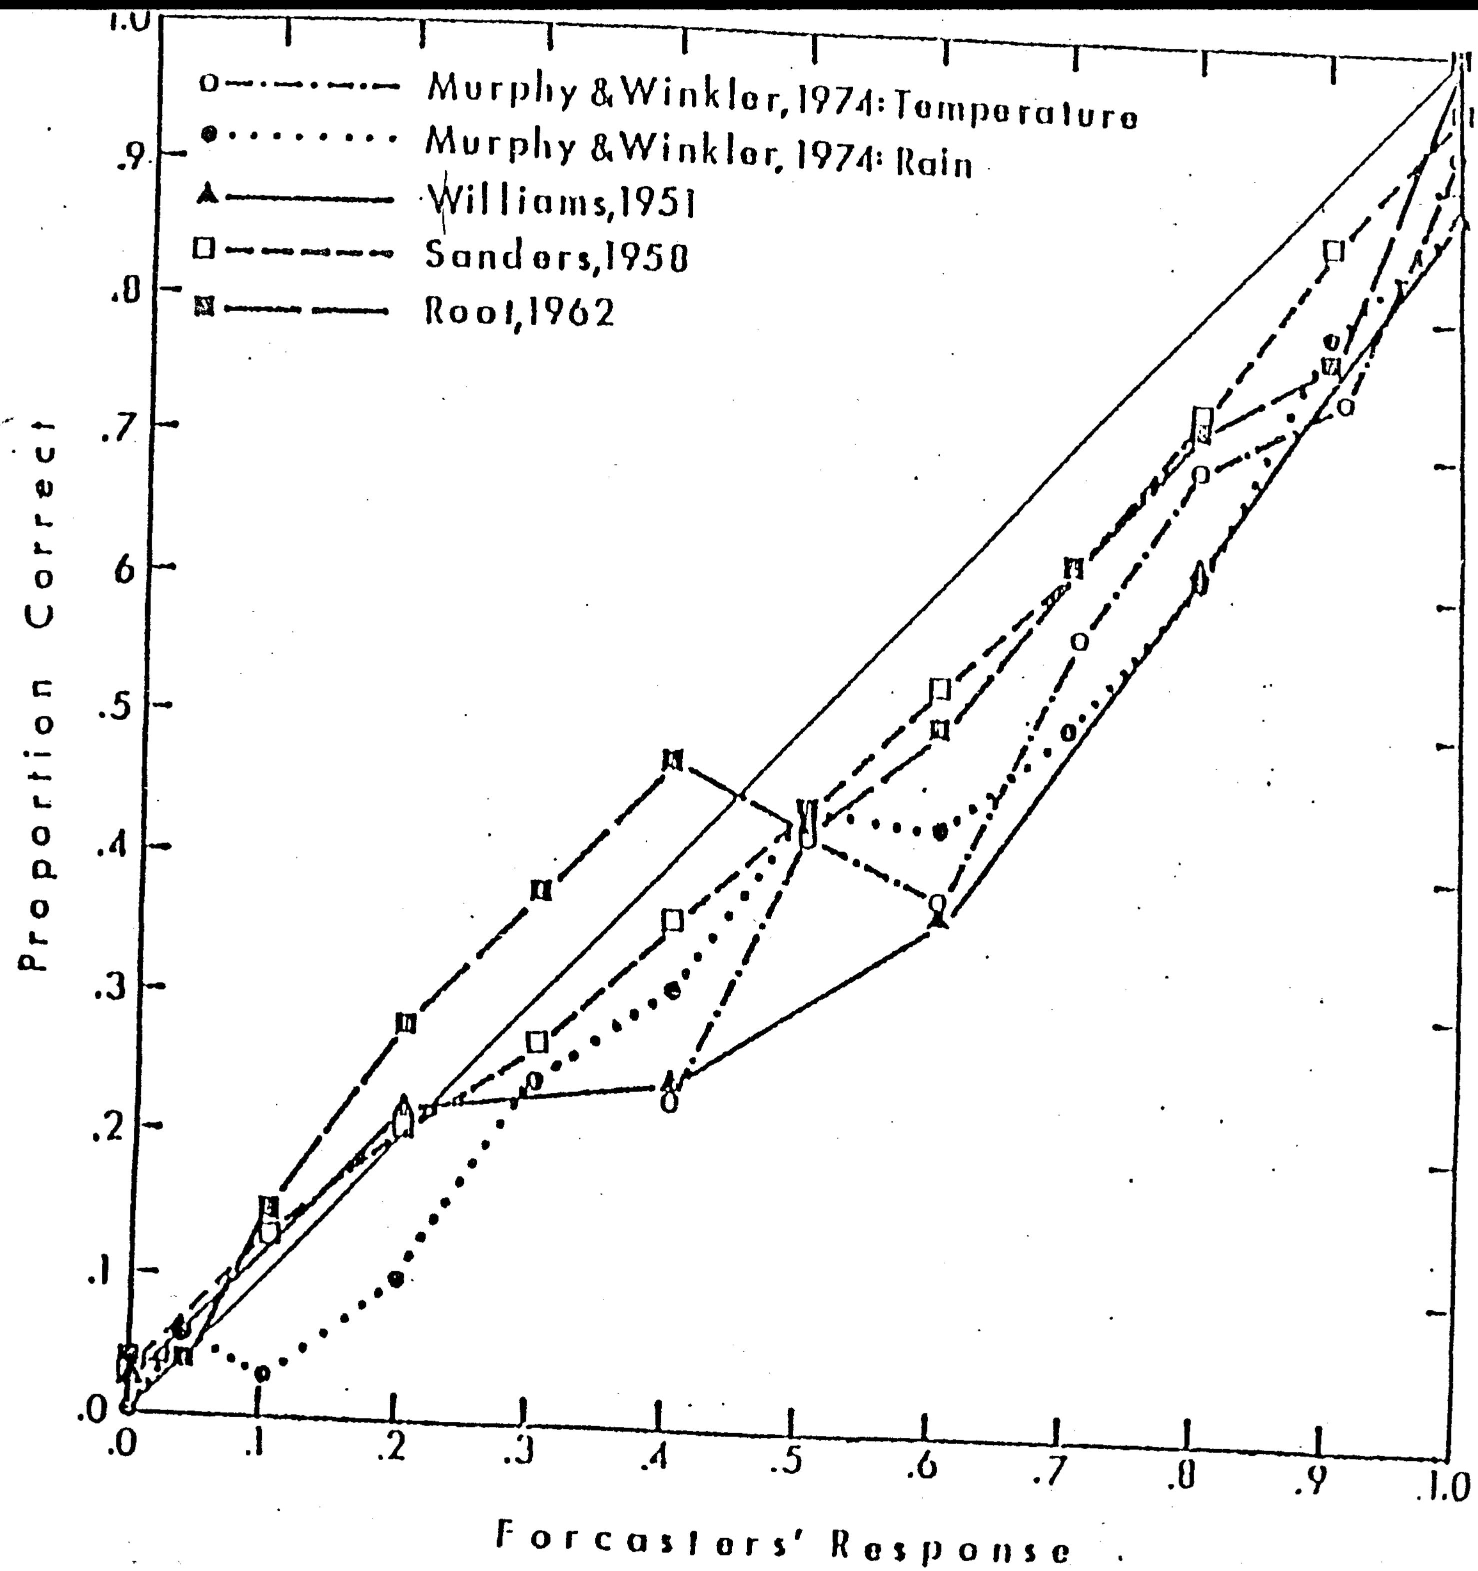
Murphy & Winkler, 1974: Temperature
Murphy & Winkler, 1974: Rain
Williams, 1951
Sanders, 1958
Root, 1962
Proportion Correct
Forcastors' Response
.0
.1
.2
.3
.4
.5
.6
.7
.8
.9
1.0

Figure 2

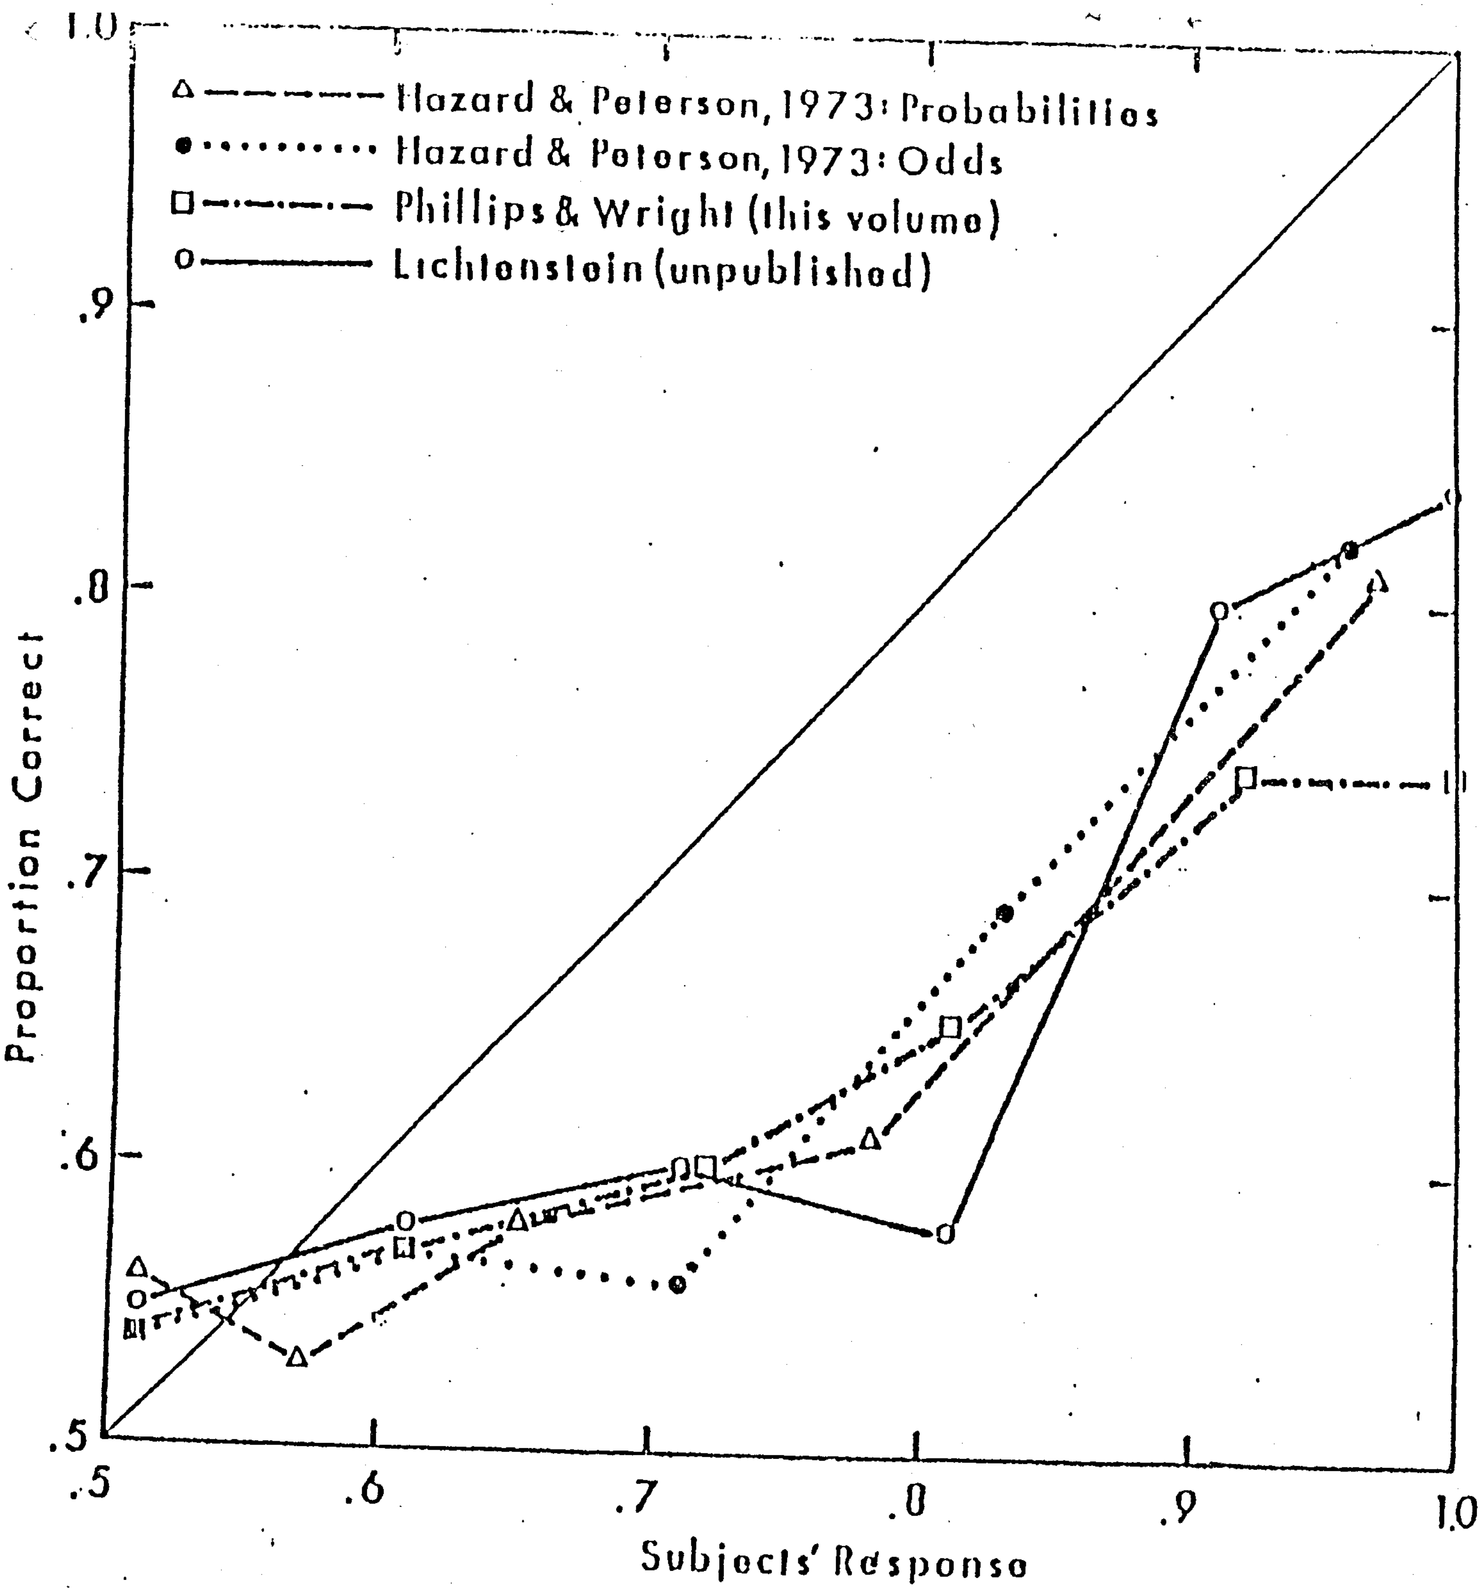
Hazard & Peterson, 1973: Probabilities
Hazard & Peterson, 1973: Odds
Phillips & Wright (this volume)
Lichtenstein (unpublished)
Proportion Correct
1.0
.9
.8
.7
.6
.5
.5
.6
.7
.8
.9
1.0
Subjects' Response

a good descriptive model of behavior, almost from the time it first came out in 1944. The most recent work in this area has been done by Kahneman and Tversky, in a paper published in Econometrica under the title of "Prospect Theory" (7).

Prospect theory is offered as an alternative to expected utility theory, and is designed as a descriptive, rather than normative, model. It is designed to predict how people make choices in uncertain situations. First of all, let me point out that expected utility theory is rarely tested in the real world, the reason being that it is very difficult to conduct such tests. To make such a test one needs objective measures of probability and enough trials to identify the underlying utility function and choice mechanism. What people who have worked in this area have done is to use either laboratory experiments or questionnaires.

One major problem with laboratory experiments is that they typically deal with only small amount of money. When a farmer has to decide which crop to plant, it is a very different choice than that faced by a student gambling for stakes of three or four dollars. There is one exception to this drawback in a paper by Binswanger (2). He has conducted experiments in India in which the stakes would have been viewed as small in the U.S., but were quite large by India standards. This appears to be a very promising method and one which might easily be used again to investigate a variety of interesting questions.

Another problem with experimental data is that they typically cannot have any subject losing money, either because of the experiment's ethics or those imposed by the funding agency. As we will see, one of the key findings in the work done by Kahneman and Tversky is that people behave differently for losses than for gains. Due to the restriction on losses their results are from hypothetical questions which mean that they must be taken with a grain of salt. We can't say too much about the reliability of the results with respect to real world applicability, but we can say that they are better than nothing. The actual results are fairly robust and do suggest that there is something going on.

Turning to Table 1, we see a summary of some of the results of this work. The notation is explained in the following way. In problem 3 (4,000, .80) indicates a gamble in which there is an 80% chance of winning 4000 and a 20% chance of winning nothing; (3000) indicates a certain payoff of 3000. The inequality sign represents preferences of the majority of subjects and the numbers in brackets appearing under each choice are the percentage of subjects choosing each gamble. The asterisk indicates significance of the results at the .01 level. So the way to read problem 3 is that 80% of the subjects preferred $3000 for sure to an 80% chance at $4000. There are several interesting results in this table.

The first thing to notice is that the left and right columns differ only by the sign of the prospect, so for example while problem 3 is (4000, .80) and (3000), problem 3' is (-4000, .80) and (-3000). You will note that the sign of the preference is reversed consistently across

Table 1*

PREFERENCES BETWEEN POSITIVE AND NEGATIVE PROSPECTS

Positive prospects				Negative prospects			
Problem 3:	(4,000, .80)	<	(3,000).	Problem 3':	(−4,000, .80)	>	(−3,000).
N = 95	[20]		[80]*	N = 95	[92]*		[8]
Problem 4:	(4,000, .20)	>	(3,000, .25).	Problem 4':	(−4,000, .20)	<	(−3,000, .25).
N = 95	[65]*		[35]	N = 95	[42]		[58]
Problem 7:	(3,000, .90)	>	(6,000, .45).	Problem 7':	(−3,000, .90)	<	(−6,000, .45).
N = 66	[86]*		[14]	N = 66	[8]		[92]*
Problem 8:	(3,000, .002)	<	(6,000, .001).	Problem 8':	(−3,000, .002)	>	(−6,000, .001).
N = 66	[27]		[73]*	N = 66	[70]*		[30]

*Kahneman, Daniel and Amos Tversky, "Prospect Theory: An Analysis of Decision under Risk," Econometrica (47), March 1979, 263-291.

columns, which indicates that risk-seeking behavior in the domain of gains is concurrent with risk-averse behavior in the domain of losses, and vice versa. This suggests that gains and losses are not treated in the same way by individuals. The next thing to notice is a comparison of problems 3 and 4, for the moment only looking at the positive side. Problem 4 is simply problem 3 with the probabilities divided by four, yet we observe a reversal in preferences. This result violates what Savage (11) called the sure-thing axiom and what some others have referred to as the strong independence axiom. Problems 7 and 8 display the same type of reversal; in this case the probabilities have been divided by 450. Another observation to be made is that in problems 3 and 7 the subjects are risk averse in the domain of gains while in problems 4 and 8 they are risk seeking. The reverse of these positions holds in each case in the domain of losses. One thing to keep in mind with respect to the next examples is that in expected utility theory prospects are evaluated in terms of their final asset position. Consider the following problems, also used by Kahneman and Tversky (7).

Problem 11: In addition to whatever you own, you have been given 1,000. You are now asked to choose between

A: (1,000, .50), and B: (500).
N = 70 (16) (84)*

Problem 12: In addition to whatever you own, you have been given 2,000. You are now asked to choose between

C: (-1,000, .50), and D: (-500).
N = 68 (69)* (31)

If one chooses A, the final asset position will be either 2,000 or 1,000, each with probability .50, while a choice of B will result in 1500 for sure. It is clear that the final asset position resulting from a choice of C is identical to A and D is identical to B. However, the results indicate a preference reversal from B to C.

To explain this reversal, Kahneman and Tversky propose a new theory which they call prospect theory. A key element in the theory is that prospects, or gambles, are valued in terms of changes in wealth position with respect to some reference point rather than in terms of final asset position. Additionally, people are assumed to use a value function to evaluate prospects. A value function is simply a utility function defined over changes rather than over final asset positions, and might appear similar to the example shown in Figure 3. Notice that the function is convex for losses and concave for gains.

Figure 3*

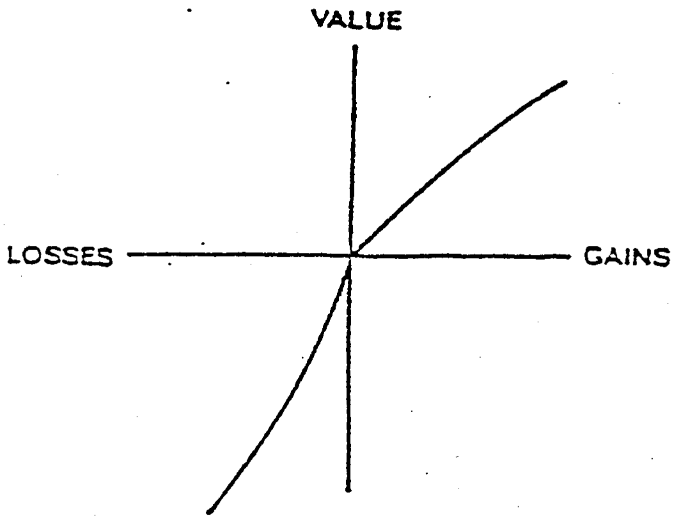

FIGURE 3.—A hypothetical value function.

Figure 4*

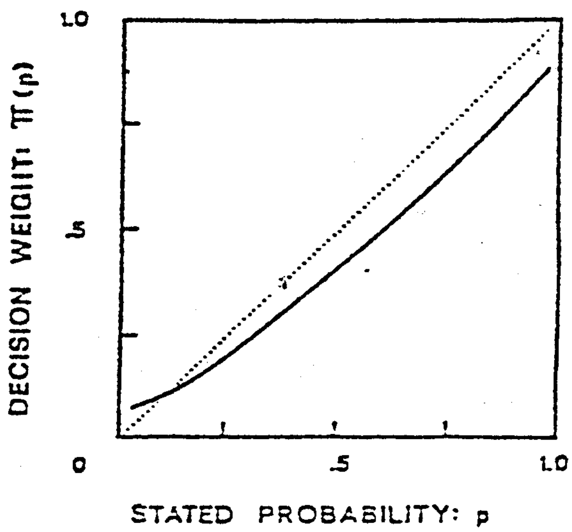

FIGURE 4.—A hypothetical weighting function.

*Kahneman, Daniel and Amos Tversky, "Prospect Theory: An Analysis of Decision under Risk, " Econometrica (47), March 1979, 263-291.

A third key concept is that losses loom greater than do gains. The intuition behind this proposition is that one can compare the feeling of having gained $100 to the feeling of having lost $100. The hypothesis is that the amount by which you feel worse for having lost $100 is greater than the amount by which you feel better for having won $100. Therefore, the loss function is steeper through the origin than is the gain function. The fourth element is a weighting function which replaces subjective probability in the decision making process. A hypothetical weighting function is pictured in Figure 4. As you can see, it is discontinuous at the end points, resulting in an overweighting of small probabilities and an underweighting of large probabilities. This can be regarded as a picture of the certainty effect discussed above, in which certain outcomes are treated differently than are nearly certain outcomes.

These are the four key elements to prospect theory, which are put together in a clear and readable fashion by Kahneman and Tversky (7). The result is a theory whereby one can predict the responses shown in Table 1. There have been other attempts at trying to develop a descriptive theory of decision making. In fact, no one part of prospect theory is completely original; each piece has appeared elsewhere in the past. Edwards (3) did work with weighting functions a long time ago and Swalm (14) observed a utility function shaped similarly to the hypothetical value function (Figure 3).

Friedman and Savage (4) also attempted to explain individual decision making under uncertainty, but their theory cannot explain all the paradoxes revealed in the Kahneman and Tversky paper (7). This is in part because Friedman and Savage's utility function is supposed to stay put; in contrast, prospect theory's value function moves around to wherever an individual's reference point is when making a particular decision.

An interesting research problem would be to come up with the equivalent of Arrow-Pratt (1, 10) measures of risk aversion for prospect theory. The Arrow-Pratt indexes are essentially measures of the degree of concavity of the utility function, i.e., how sharply bowed the function is. The value function is convex for losses and concave for gains, so you cannot really use a single measure. It would be interesting to try to develop a measure that would be equivalent to risk aversion but would apply to prospect theory. I think this measure will end up being some part of "loss aversion." For example, suppose you offer someone a gamble wherein a coin is flipped, with heads being a win of $150 and tails being a loss of $100. Suppose the person turns you down. If you ask an economist to explain why they turned it down, the reply most likely will be risk aversion. That is hogwash for the plain and simple reason that $100-$150 is small relative to a typical person's wealth. It is unlikely that there is enough curvature in the utility function over that small a range to give you such a result. For a typical person, the curve over that range will look pretty close to linear. Intuitively, the reason is that people do not want to lose $100, and maybe they do not want to have to go home to tell their families that they have lost $100. I think that a much better measure for explaining this kind of behavior would be some kind of loss aversion measure. This would be a measure of how much steeper the loss function is relative to the gain function. It might be interesting in a farm context to try and estimate some of these loss aversion measures and see whether they do a better job of explaining behavior, relative to risk aversion measures.

Other than the value function/reference point contribution, the other major lesson to be learned from prospect theory and related work is what economists might refer to as the structure or form of a problem. The way the problem is stated appears to make a big difference. Problems like 11 and 12, or the two versions of the taxi cab problem, elicit different responses, although they are identical in substance, because the form of presentation is different. In a sense prospect theory is a theory of form. It tells you that the manner in which a problem is presented makes a difference. The moral from this is that if you want farmers, or any other people, to respond in a particular way, it can help you to determine the way in which the problem should be explained.

Most of the applications I have been thinking about are in the domain of marketing. One way of thinking about this is a very sophisticated version of the famous book How to Lie with Statistics (6). How can data be presented factually but in a way so that people are likely to behave the way you want them to? It could be because you want them to buy insurance, or plant a certain crop, or whatever. It is clear that you can present probabilistic information to people in a lot of different ways, and what all of this research should tell you is that what they will do will depend as much, if not more, on the way in which you present that information as on the factual content of the information itself.

References

(1) Arrow, K. J., Essays in the Theory of Risk-Bearing. Chicago: Markham, 1971.

(2) Binswanger, H. P. "Attitudes towards Risk: Experimental Measurement in Rural India." Center Discussion Paper No. 285, Economic Growth Center, Yale University, June 1978.

(3) Edwards, W. "Subjective Probabilities Inferred from Decisions." Psychological Review (69), 1962, 109-135.

(4) Friedman, M., and L. J. Savage. "The Utility Analysis of Choices Involving Risks." Journal of Political Economy (56), 1948, 279-304.

(5) Grether, D. M. "Bayes Rule as a Descriptive Model: The Representativeness Heuristic." Social Science Working Paper 245, California Institute of Technology, January 1979.

(6) Huff, D. How to Lie with Statistics. New York: Norton, 1954.

(7) Kahneman, D., and A. Tversky. "Prospect Theory: An Analysis of Decision under Risk." Econometrica (47), March 1979, 263-291.

(8) Lichtenstein, S. Slovic, P., Fischhoff, B. Layman, M., and B. Combs. "Judged Frequency of Lethal Events." Journal of Experimental Psychology: Human Learning and Memory (4), 1978, 551-578.

(9) Lichtenstein, S., Fischhoff, B., and L. D. Phillips. "Calibration of Probabilities: The State of the Art," in Decision Making and Change in Human Affairs, H. Jungermann and G. de Zeeuw (eds.), 1977.

(10) Pratt, J. W. "Risk Aversion in the Small and in the Large." Econometrica (32), 1964, 122-136.

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

(12) Shefrin, H. M., and R. Thaler. "Rules and Discretion in a Two-Self Model of Intertemporal Choice," Cornell University mimeo, May 1980.

(13) Slovic, P., Fischhoff, B., and S. Lichtenstein. "Behavioral Decision Theory." Annual Review of Psychology (28), 1977, 1-39.

(14) Swalm, R. O. "Utility Theory - Insights into Risk Taking." Harvard Business Review (44), 1966, 123-136.

(15) Thaler, R. "Individual Intertemporal Choice: A Preliminary Investigation," Research Memorandum, Cornell University, February 1979.

(16) Thaler, R. "Toward a Positive Theory of Consumer Choice." Journal of Economic Behavior and Organization (1), 1980, in press.

(17) Thaler, R., and H. M. Shefrin. "An Economic Theory of Self-Control." Forthcoming in Journal of Political Economy.

(18) Tversky, A., and D. Kahneman. "Judgment under Uncertainty: Heuristics and Biases." Science (185), 1974, 1124-1131.

(19) Tversky, A., and D. Kahneman. "Causal Schemata in Judgments under Uncertainty," in Progress in Social Psychology, M. Fishbein (ed.), Hillsdale: Lawrence Erlbaum Associates, 1977.

(20) Von Neumann, J., and O. Morgenstern. Theory of Games and Economic Behavior. Princeton: Princeton University Press, 1944.