

PARADOXES AND VIOLATIONS OF NORMATIVE DECISION THEORY

JAY SIMON
Defense Resources Management
Institute, Naval Postgraduate
School, Monterey, California

YITONG WANG
L. ROBIN KELLER
The Paul Merage School of
Business, University of
California Irvine, Irvine,
California

Traditional economic decision theory proposes that people behave in certain ways when faced with a well-formulated set of alternatives and information. It is a normative theory; it suggests that people should act according to certain decision rules, but not that they will necessarily do so in reality. The standard assumption is that a person has a utility function over all possible outcomes, which defines the desirability or usefulness of any given scenario. Traditional utility theory asserts that in a decision without uncertainty, a person should choose the alternative resulting in the highest level of utility, and that in a decision with uncertainty, a person should choose the alternative with the highest expected utility. Detailed explanations are given by von Neumann and Morgenstern and Savage [1,2].

These ideas are generally sound and can be valuable when used as a guide for decision makers. However, it was initially thought that they were also valuable as a descriptive theory. It has since been determined that people violate normative decision theory in many clear and predictable ways. In this article, we introduce and discuss several of these interesting behavioral phenomena.

One of the earliest researchers to demonstrate a deviation from expected utility theory was Allais [3]. He distributed two surveys. In the first survey, respondents were asked to choose between the following

two options (original amounts were in francs):

Choice 1.

A:	\$1 million for sure
B:	0.10 probability of \$2 million 0.89 probability of \$1 million 0.01 probability of \$0

In the second survey, respondents were asked to choose between the following two options:

Choice 2.

C:	0.11 probability of \$1 million 0.89 probability of \$0
D:	0.10 probability of \$2 million 0.90 probability of \$0

The majority of respondents preferred *A* to *B*, and preferred *D* to *C*.

Choosing *A* and *D* appears to be reasonable. However, this example has been called the *Allais Paradox*, because these two choices contradict the idea that people are maximizing expected utility. If we let U_0 represent the person's utility for \$0, U_1 represent the person's utility for \$1 million, and U_2 represent the person's utility for \$2 million, then preferring *A* to *B* would imply that

$$U_1 > 0.10U_2 + 0.89U_1 + 0.01U_0,$$

or equivalently,

$$U_1 > \frac{10}{11}U_2 + \frac{1}{11}U_0.$$

If we examine the second survey choice, we find that preferring *D* to *C* implies that

$$0.11U_1 + 0.89U_0 < 0.10U_2 + 0.90U_0.$$

This may not seem like a problem at first glance. However, with simple

algebraic manipulation, we see that this is equivalent to

$$U_1 < \frac{10}{11}U_2 + \frac{1}{11}U_0.$$

This is the exact opposite result of what we saw in the first survey! It means that regardless of the utilities of the three outcomes, choosing both *A* and *D* (or both *B* and *C*) contradicts the theory of expected utility maximization.

Table 1 reveals the structure of the transformation of the first choice between *A* and *B* into the second choice between *C* and *D*. The final column corresponds to the 0.89 chance of the \$1 million outcome received in either *A* or *B* in the first choice. This \$1 million was replaced with a \$0 in both *C* and *D* in the second choice. Changing the amount of this “sure thing” of getting the same outcome of \$1 million in the 0.89 probability outcome in choice 1 to a different sure thing of \$0 in choice 2 should not change the preference order between the options. Savage (2, Postulate 2) called this the *sure-thing* principle, which must be obeyed under expected utility maximization.

This was one of the first clear demonstrations of a systematic violation of the principles of expected utility theory. That in itself is a very influential finding. However, we may want to ask why it happened. What element of people’s underlying preferences is not being captured by expected utility theory?

It turns out there are multiple explanations for Allais’ results, but the main reason they occurred is a phenomenon often referred to as the *certainty effect*. Reducing the probability of a negative outcome from 0.01 to 0 is usually judged to be a greater improvement than reducing the probability of that same negative outcome from, say, 0.90 to 0.89. Expected utility maximization requires

Table 1. Allais Paradox

	0.10	0.01	0.89
A	\$1 million	\$1 million	\$1 million
B	\$2 million	\$0	\$1 million
C	\$1 million	\$1 million	\$0
D	\$2 million	\$0	\$0

that preferences over gambles are linear in probabilities, and in Allais’ example, they clearly are not.

Another major finding in this area was discussed by Ellsberg [4]. He conducted an experiment in which subjects were presented with a jar. The subjects were told that the jar contains 30 red balls, and 60 balls which are some unknown mix of yellow and black. They were asked to choose between two gambles:

-
- A: Receive \$100 if the ball is red
 - B: Receive \$100 if the ball is black
-

The majority of participants chose gamble *A*. They were then asked to choose between the following two gambles:

-
- C: Receive \$100 if the ball is red or yellow
 - D: Receive \$100 if the ball is black or yellow
-

Given this choice, the majority of participants selected gamble *D*. As in the Allais’ example, choosing both *A* and *D* seems reasonable. However, this has been called the *Ellsberg Paradox*, because it is also a violation of expected utility theory. It may not be immediately clear how expected utility is violated, so let us write a participant’s estimated fraction of black balls as *B*, and observe the following contradiction:

Choosing *A* implies that

$$\frac{1}{3}U(\$100) + \frac{2}{3}U(\$0) > BU(\$100) + (1 - B)U(\$0).$$

Assuming that $U(\$100) > U(\$0)$, this simplifies to $B < 1/3$.

Choosing *D* implies that

$$(1 - B)U(\$100) + BU(\$0) < \frac{2}{3}U(\$100) + \frac{1}{3}U(\$0).$$

Assuming that $U(\$100) > U(\$0)$, this simplifies to $B > 1/3$.

Thus, there is no consistent set of beliefs (about the fraction of black balls and the fraction of yellow balls) which allows us, using traditional subjective expected utility theory, to prefer both gamble *A* and gamble *D*. That is, no single probability of drawing a black ball would lead to this pair of responses. However, there is a simple reason that most people choose both gamble *A* and gamble *D*: these are the choices for which we can easily compute the exact probability of receiving \$100. Ellsberg showed that not only do people display aversion to risk, they also display aversion to “ambiguity.” People tend to prefer gambles for which they are confident that they know the exact probabilities involved.

Table 2 shows the structure of this paradox. The choice between *A* and *B* is transformed by changing the sure thing of \$0 when a yellow ball is chosen to a sure thing of \$100 when yellow is chosen. So, if *A* is preferred over *B*, then *C* should be preferred over *D*.

Possibly the most groundbreaking work to bring violations of normative utility theory to a wide audience is from Kahneman and Tversky [5]. They highlighted two major behavioral violations of expected utility maximization, and demonstrated them using simple survey results.

Their first major violation resulted from a very simple pair of questions:

1. Would you prefer to receive \$3000 for sure, or \$4000 with probability 0.8?
2. Would you prefer to lose \$3000 for sure, or lose \$4000 with probability 0.8?

Expected utility maximization does not tell us precisely what the answers to these questions should be. However, notice that in

the second question, the expected value of losing \$4000 with probability 0.8 is $-\$3200$. It has greater risk than the sure loss, and a lower expected value. Preferring the gamble would imply that the person is risk-seeking. Utility curves formed over total wealth are almost always concave, implying risk aversion.

The answers to these two questions were startling. People (in aggregate) overwhelmingly preferred receiving \$3000 for sure in question 1, and overwhelmingly preferred losing \$4000 with probability 0.8 in question 2. That is, they were clearly risk-averse in question 1, and clearly risk-seeking in question 2. Kahneman and Tversky also asked similar types of questions in different contexts, using different probabilities and outcomes. In all cases, they found that the majority’s preferred answer to the question dealing with losses was the opposite of that for the question dealing with gains. They referred to this phenomenon as the *reflection effect*.

Since these people did not all have the same level of wealth, it would be impossible to construct a utility function over wealth that could incorporate the reflection effect and accurately represent these subjects’ responses. Instead, Kahneman and Tversky proposed that people evaluate outcomes as gains and losses rather than using the resulting total wealth levels. Based on their survey data, they postulated that people tend to be risk-averse over gains and risk-seeking over losses. Numerous subsequent studies have demonstrated this phenomenon, yet it is still not fully accepted or incorporated into traditional economic theory.

Kahneman and Tversky also found nonlinearities in probability, and expressed them in some more detail than Allais. In addition to certainty effects, they found that people tend to place too much weight on outcomes with very small (nonzero) probabilities. That is, they do not fully understand or incorporate how truly unlikely these outcomes are. They postulated that people treat probabilities according to the weighting function shown in Fig. 1. The horizontal axis represents the actual given probability, and the vertical axis represents the probability that people implicitly interpret it to be

Table 2. Ellsberg Paradox

	30 Balls	60 Balls are Either Black or Yellow	
	Red balls	Black balls	Yellow balls
A	\$100	\$0	\$0
B	\$0	\$100	\$0
C	\$100	\$0	\$100
D	\$0	\$100	\$100

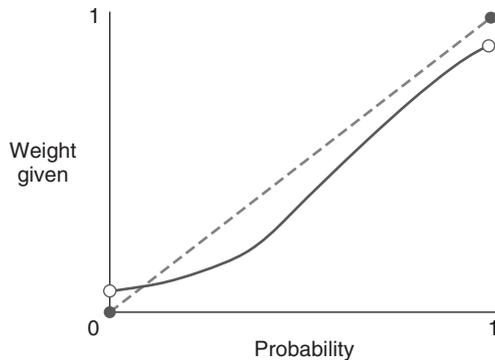


Figure 1. Nonlinear application of probabilities as observed by Kahneman and Tversky [5].

when making decisions. The “correct” interpretation of probability, as implied by traditional expected utility theory, is given by the 45 degree dotted line. This graph incorporates the certainty effect (at both zero and one), the overweighting of small probabilities, and the underweighting of middle and higher probabilities.

Another related behavioral phenomenon is what Tversky and Kahneman [6] referred to as *loss aversion*. Loss aversion states that a loss is deemed to be more significant than an equal-sized gain. The theoretical loss aversion concept explains two other more specific observed concepts. These two specific concepts are the *endowment effect* and the *status quo bias*.

The *endowment effect* is discussed by Thaler [7]. The premise is that the loss in utility from giving up an item is greater than the gain in utility from acquiring it. It was demonstrated experimentally by Kahneman *et al.* [8]. They asked subjects to place a monetary value on a cheap decorative mug. The catch was that some subjects had been given the mug beforehand, and others had not. Since the mugs were assigned randomly, there was no apparent reason that the underlying preferences should differ between those with mugs and those without. However, they found that the subjects who were given the mug placed *much* higher values on it than those who were not. In one experiment, the median value for the mug was \$7.12 among the subjects who had one, and \$3.12 among those who had not. In a second experiment,

these median values were \$7 and \$3.50, respectively. This was clear evidence that people tend to place much higher values on items they possess, or are “endowed” with.

The *status quo bias* was discussed by Samuelson and Zeckhauser [9]. *Status quo bias* is a preference for remaining in the current situation. The example that the authors used dealt with financial investments. They conducted an experiment in which several different funds were described to the participant. When participants were told that they currently had money invested in one of the funds, then they generally preferred that fund to the others. When they were given a “neutral” scenario in which they did not have money invested in any of the funds, then no such bias was observed.

Loss aversion explains both the endowment effect and the status quo bias. If a loss is more significant than an equal-sized gain (in terms of utility), then losing an item is clearly more significant than receiving that same item. It is not difficult to use loss aversion to explain the endowment effect. To understand the status quo bias, we can interpret a change as “losing” the current scenario and obtaining the new one. If the two scenarios are equally desirable when viewed from a neutral perspective, then loss aversion suggests that the decision maker will prefer the status quo.

Notice that in each of the two preceding examples, if the outcomes were viewed in terms of raw overall outcomes rather than gains and losses, the groups would be facing identical questions. In the endowment effect example, both groups would be expressing the difference in monetary value between a “mug” outcome and a “no mug” outcome. In the status quo bias example, each individual would be simply selecting the most desirable fund.

One implication of expected utility theory is that a person’s chosen course of action should not depend on the manner in which the alternatives are presented. If two decision situations involve the same outcomes with the same probabilities resulting from the corresponding alternatives, then the person’s decision should be identical for both. This property is called *invariance*. It is discussed by von Neumann and Morgenstern [1],

and studied further by Arrow [10]. Invariance is often violated in reality. These violations are referred to as *framing* effects.

Tversky and Kahneman [11] conducted a pair of surveys which demonstrated a clear framing effect. In both survey scenarios, a group of 600 people is about to be exposed to a deadly disease. In the first survey, the choice is whether to save 200 people for sure, or save all 600 with probability 1/3 (and save nobody with probability 2/3). The vast majority of participants chose the sure option. In the second survey, the choice is whether to allow 400 people to die for sure, or accept a 2/3 probability of all 600 people dying (and a 1/3 probability of nobody dying). The vast majority of participants chose the gamble.

Look at those two survey questions again. They are actually asking the same thing! One is framed in terms of number of lives saved, and the other is framed in terms of number of deaths. The outcomes and their associated probabilities, however, are identical. Simply due to the ways in which the situations are described, we find that many people's preferences are completely reversed. An observant reader may notice that this particular framing effect occurs because one question is described in terms of gains, and the other in terms of losses. Tversky and Kahneman [12] discussed this in greater detail. As prospect theory tells us, people tend to be risk-averse for gains and risk-seeking for losses.

One of the arguments in favor of using a risk-averse utility curve over wealth is the observation that most people will never accept a 50/50 gamble between gaining and losing $\$X$, regardless of the value of X . This implies a concave (risk-averse) utility function for wealth. Loss aversion is an alternative explanation for this observation; it asserts that " $-\$X$ " has a larger effect on utility than " $+\$X$ " does, without having to assess or incorporate wealth level.

Loss aversion, along with Kahneman and Tversky's earlier data, implies that rather than looking at potential overall wealth levels when making decisions, people tend to view outcomes as gains and losses relative to a particular reference level. Kahneman and Tversky's [5] prospect theory suggests that

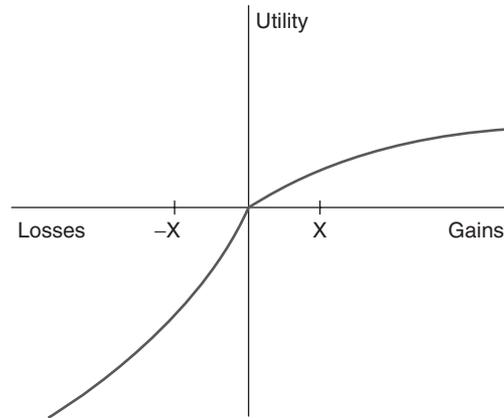


Figure 2. The utility function associated with prospect theory Kahneman and Tversky [5].

people act according to a utility function¹ similar to the one shown in Fig. 2, where the slope of the utility is steeper when a loss of $\$X$ is incurred than when a gain of $\$X$ is incurred. This idea alone can explain many real-world deviations from traditional expected utility theory. Baucells and Rata [13] demonstrated empirically that using a reference point in this manner is an important factor in risk-taking behavior; Camerer [14] provides many additional real-world examples.

Given that people do not always conform to expected utility, there are different ways to proceed. One is to investigate how to aid people to do so. For example, Keller [15,16] investigated how different visual representations of gambles (most promisingly drawings of color coded balls labeled with the monetary amount that would be received if a ball is drawn) could lower violations of expected utility for questions similar to the Allais Paradox. Another approach is to relax the assumptions required by expected utility and

¹Kahneman and Tversky refer to this as a value function rather than a utility function, and it is generally understood to be an appropriate model for preferences under certainty. However, they proposed it initially based on observed preferences over gambles, as we have discussed earlier. Thus, to be consistent with our usage of the terms "value" and "utility," we refer to the curve in Fig. 2 as a utility function.

construct generalizations of expected utility, such as models that relax the requirement of linearity in probabilities. Kahneman and Tversky's [5] prospect theory is the most well-known generalized utility theory, but there are many others.

Becker and Sarin [17] introduced a lottery-dependent utility model that allows the utility achieved in an outcome to vary according to the lottery in which it occurs, but requires the true probabilities to be used. Their model was rather general, but can be used to explain some violations of expected utility. Daniels and Keller [18] tested this model experimentally, and found that it was useful in predicting choice among gambles, but not noticeably better than expected utility when actually determining decision makers' underlying preferences. Some other generalized utility models include weighted utility [19], regret theory (20, 21), and skew-symmetric bilinear utility [22,23]. For more examples and detailed discussion the reader is directed to Refs 24, 25, or 26.

In general, expected utility theory is quite valuable as a normative theory for guiding good decision making. However, between the certainty effect, ambiguity aversion, the reflection effect, the endowment effect, status quo bias, loss aversion, and framing effects, it is painfully apparent that people violate the traditional theory in many clear and predictable ways.

A final question for thought: in the real world, how often do these sorts of paradoxes and violations occur? Certainly the effects and biases that we have discussed do exist, but how often do they lead to incorrect or suboptimal decisions? If they occur very infrequently, then they are interesting as academic issues only. However, if they do come up reasonably often in the real world, then they undoubtedly merit further study and awareness.

REFERENCES

1. von Neumann J, Morgenstern O. *Theory of games and economic behavior*. Princeton (NJ): Princeton University Press; 1947, second edition.
2. Savage L. *The foundations of statistics*. New York: John Wiley & Sons, Inc.; 1954.
3. Allais M. Le comportement de l'homme rationnel devant le risque: critique des postulats et axiomes de l'Ecole Americaine. *Econometrica* 1953;21(4):503–546.
4. Ellsberg D. Risk, ambiguity, and the Savage axioms. *Q J Econ* 1961;75(4):643–669.
5. Kahneman D, Tversky A. An analysis of decision under risk. *Econometrica* 1979; 47(2):263–292.
6. Tversky A, Kahneman D. Loss aversion in riskless choice: a reference dependent model. *Q J Econ* 1991;107(4):1039–1061.
7. Thaler R. Toward a positive theory of consumer choice. *J Econ Behav Organ* 1980; 1(1):39–60.
8. Kahneman D, Knetsch JL, Thaler RH. Experimental tests of the endowment effect and the Coase theorem. *J Polit Econ* 1990; 98(6):1325–1348.
9. Samuelson W, Zeckhauser R. Status quo bias in decision making. *J Risk Uncertain* 1988;1(1):7–59.
10. Arrow K. Risk perception in psychology and economics. *Econ Inq* 1982;20(1):1–9.
11. Tversky A, Kahneman D. The framing of decisions and the psychology of choice. *Science* 1981;211(4481):453–458.
12. Tversky A, Kahneman D. Rational choice and the framing of decisions. *J Bus* 1986;59(4):S251–S278.
13. Baucells M, Rata C. A survey study of factors influencing risk-taking behavior in real-world decisions under uncertainty. *Decis Anal* 2006;3:163–176.
14. Camerer C. Prospect theory in the wild: evidence from the field. In: Kahneman D, Tversky A, editors. *Choices, values, and frames*. Cambridge: Cambridge University Press; 2001. pp. 288–300.
15. Keller LR. The effects of problem representation on the sure-thing and substitution principles. *Manage Sci* 1985a;31(6):738–751.
16. Keller LR. Testing the “reduction of compound alternatives” principle. *Omega, Int J Manage Sci* 1985b;13(4):349–358.
17. Becker J, Sarin R. Lottery dependent utility. *Manage Sci* 1987;33:1367–1382.
18. Daniels RL, Keller LR. An experimental evaluation of the descriptive validity of lottery-dependent utility theory. *J Risk Uncertain* 1990;3(2):115–134.
19. Chew SH, Waller WS. Empirical tests of weighted utility theory. *J Math Psychol* 1986;30:55–72.

20. Bell D. Regret in decision making under uncertainty. *Oper Res* 1982;30:961–981.
21. Loomes G, Sugden R. Regret theory: an alternative theory of rational choice under uncertainty. *Econ J Nepal* 1982;92:805–824.
22. Fishburn PC. Transitive measurable utility. *J Econ Theory* 1983;31:293–317.
23. Fishburn PC. SSB utility theory: an economic perspective. *Math Soc Sci* 1984;8:63–94.
24. Edwards W, editor. *Utility theories: measurements and applications*. Boston (MA): Kluwer; 1992.
25. Fishburn PC, LaValle I, editors. *Annals of operations research 19: choice under uncertainty*. Basel: J. C. Baltzer; 1989.
26. Weber M, Camerer C. Recent developments in modelling preferences under risk. *OR Spektrum* 1987;9:129–151.