



BA 513/STA 234: Ph.D. Seminar on Choice Theory
Professor Robert Nau
Spring Semester 2008

Notes and readings for class #4: Non-expected utility theory (revised February 6, 2008)

Primary readings:

- 1a. "Alternatives To Expected Utility: Foundations" by Robert Sugden. In *Handbook of Utility Theory Vol. 2 Extensions* edited by Salvador Barberà, Peter Hammond, and Christian Seidl
- 1b. "Recent Tests Of Generalizations Of Expected Utility Theory" by Colin Camerer, in *Utility Theories: Measurements and Applications* edited by Ward Edwards, 1992

Supplementary readings:

- 2a. "Prospect theory: an analysis of decision under risk" by Daniel Kahneman and Amos Tversky, *Econometrica*, 1979, v. 47, no. 2, pp. 263-289
- 2b. "'Expected Utility' Analysis without the Independence Axiom" by Mark Machina *Econometrica*, 1982, v. 50, no. 2, pp. 277-323
- 2c. "Choice under uncertainty: problems solved and unsolved" by Mark Machina, *Journal of Economic Perspectives* 1987
- 2d. "The Dual Theory of Choice under Risk" by Menachem Yaari, *Econometrica* 1987, v. 55, no. 1, pp. 95-115
- 2e. "Subjective Probability and Expected Utility without Additivity" by David Schmeidler *Econometrica*, 1989, v. 57, no. 3, pp. 571-587
- 2f. "Maxmin Expected Utility with Non-unique Prior" by Itzhak Gilboa and David Schmeidler, *J. Math. Econ.* 1989, v. 18, pp. 141-153
- 2g. "Dynamic consistency and non-expected utility models of choice under uncertainty" by Mark Machina, *Journal of Economic Literature* 1989
- 2h. "A More Robust Definition of Subjective Probability" by Mark Machina and David Schmeidler, *Econometrica* 1992

- 2i. “Advances in Prospect Theory: Cumulative Representation of Uncertainty” by Amos Tversky and Daniel Kahneman, *Journal of Risk and Uncertainty* 1992
- 2j. “On The Intuition Of Rank-Dependent Expected Utility” by Enrico Diecidue and Peter Wakker, *Journal of Risk and Uncertainty*, 2001
- 2j. “An Axiomatization of Cumulative Prospect Theory” by Peter Wakker and Amos Tversky, *Journal of Risk and Uncertainty*, 1993
- 2k. “A Smooth Model of Decision Making Under Ambiguity” by Peter Klibanoff, Massimo Marinacci, and Sujoy Mukerji, *Econometric*, 2005,
<http://www.kellogg.northwestern.edu/faculty/klibanof/ftp/workpap/smoothambiguity.pdf>
- 2l. “Uncertainty Aversion with Second-Order Utilities and Probabilities” by Robert Nau, *Management Science* 2006, <http://faculty.fuqua.duke.edu/~rnau/soup.pdf>

The non-expected utility revolution

The 1960’s and 1970’s were a period in which the by-then-standard paradigm of expected-utility and subjective-expected-utility theory was refined and applied to problems of decision analysis and economic modeling. By the late 1970’s, however, dissatisfaction with EU and SEU theory began to emerge on a number of fronts: behavioral decision theorists and experimental economists began to uncover new evidence that subjects systematically violate the axioms of utility theory in a variety of settings, and microeconomists and some decision analysts began to feel that the theory placed too-severe constraints on the modeling of rational behavior. Many different patterns of non-EU behavior have been found in laboratory experiments since then, as discussed in the article by Camerer, but much of the deep theory in this area has been motivated by two simple thought-experiments, the famous “paradoxes” constructed by Allais (1953) and Ellsberg (1961).

Allais’ paradox involves a comparison between two pairs of acts that yield potentially large amounts of money in three states of the world. The states could be outputs of a random device such as a roulette wheel (in which case their probabilities would be exactly known) or they could be naturally-occurring events whose probability is only approximately known. The acts and their probabilities (exact or approximate) are shown in the table below. Thus, for example, act f yields \$500K for sure, while act g yields a 1% chance of getting nothing, a 10% chance of getting \$2.5M, and an 89% chance of getting \$500K.

	$p \approx .01$	$p \approx .10$	$p \approx .89$
f	\$500K	\$500K	\$500K
g	\$0	\$2,500K	\$500K
f'	\$500K	\$500K	\$0
g'	\$0	\$2,500K	\$0

If asked to choose between f and g , most individuals prefer to take the sure thing (f) precisely because it is a sure thing: they are not willing to risk even a small (1%) chance of getting nothing in order to increase the payoff from \$500K to \$2.5M in another state that has only 10% probability. But if they are asked to choose between f' and g' , neither of which is a sure thing, most prefer g' because it yields a chance at a really huge payoff that is only slightly less likely than the chance of the much smaller, though still large, payoff (a 10% chance of a jackpot of \$2.5M rather than an 11% chance of “only” \$500K). Ask yourself if you would make the same choices! This pattern violates the independence axiom of both EU and SEU theory, because the independence axiom says that preferences between acts that agree in some states (the most-likely state in this case) should not depend on *how* they agree there. This violation of the independence axiom is usually interpreted to show that the utility of an act is not always a linear function of the probabilities of the states, particularly when probabilities close to 1 or 0 are involved. Here, it appears that the difference between a 99% and 100% chance of winning a really big prize is more psychologically important than the exact size of the prize, while the difference between a 10% chance and an 11% chance of winning is *less* important than the size of the prize.

Ellsberg’s paradox demonstrates a violation of SEU rather than EU theory, because it directly addresses the question of whether there is a difference between “objective” and “subjective” uncertainty. The experiment involves comparison of acts in which the receipt of a prize is contingent on the color of balls drawn from urns, conjuring up an image of a TV game show. There are two versions, a three-color version and a two color version. The setup for the 2-color problem is as follows:

Urn 1 contains 50 red balls and 50 black balls.

Urn 2 contains 100 red and black balls in unknown proportions.

A single ball will be drawn from each urn.

Which of these two acts would you prefer:

f = “win \$100 if the ball drawn from urn 1 is red”

g = “win \$100 if the ball drawn from urn 2 is red”

Now which of *these* two acts would you prefer:

f' = “win \$100 if the ball drawn from urn 1 is black”

g' = “win \$100 if the ball drawn from urn 2 is black”

	50 balls		50 balls	
Urn 1 (known)	<i>red</i>	<i>red</i>	<i>black</i>	<i>black</i>
Urn 2 (unknown)	<i>red</i>	<i>black</i>	<i>red</i>	<i>black</i>
<i>f</i>	\$100	\$100	\$0	\$0
<i>g</i>	\$100	\$0	\$100	\$0
<i>f'</i>	\$0	\$0	\$100	\$100
<i>g'</i>	\$0	\$100	\$0	\$100

Most persons strictly prefer to bet on urn 1 if the winning color is red ($f > g$) and they *also* strictly prefer to bet on urn 1 if the winning color is black ($f' > g'$), on grounds that they prefer the winning event to have an objective probability of 50% rather than a subjective probability that could be anywhere between 0 and 1, even if they have no reason to believe that one color is more likely than the other to be drawn from the second urn. Now, f and g agree with each other in both the red-red state and the black-black state, and so do f' and g' , so the independence axiom requires that the choice in both cases can only depend on what happens in the red-black and black-red states. But f is the same as g' in these remaining states, and g is the same as f' there, so $f > g$ must imply $g' > f'$. More transparently, $f > g$ suggests that the subject believes that there are more black than red balls in urn 2, while $f' > g'$ suggests exactly the opposite.

	30 balls		60 balls	
	<i>red</i>	<i>black</i>	<i>yellow</i>	
<i>f</i>	\$0	\$100	\$100	
<i>g</i>	\$100	\$0	\$100	
<i>f'</i>	\$0	\$100	\$0	
<i>g'</i>	\$100	\$0	\$0	

In Ellsberg's 3-color paradox a single urn contains 90 balls of (potentially) three colors: 30 are known to be red while the other 60 are black and yellow in unknown proportions. In a comparison of acts f and g in the table above, most subjects prefer f because it yields a known probability of 2/3 of winning the \$100 prize rather than an ambiguous probability that could be anywhere between 1/3 and 1, depending on the proportions of black and yellow. But in a

comparison of f' and g' they prefer g' because it yields a known probability of 1/3 of winning rather than an ambiguous probability that could be between 0 and 2/3. This behavior is an even more obvious violation of independence, because the two pairs of acts differ only in that an agreeing payoff of \$0 in the yellow state has been substituted for an agreeing payoff of \$100.

These two paradoxes were largely dismissed as curiosities until the 1970's, when behavioral researchers began to study choice under uncertainty in more carefully controlled experiments¹, and economic theorists started to question the authority of the vNM and Savage axioms that had previously been taken to be self-evident. A host of “non-expected utility” models of preferences suddenly exploded on the scene in the early 1980's. Those were heady days for decision theory: an earthquake had shaken the field and a paradigm shift seemed imminent. Disciplinary boundaries became blurred as economists, psychologists, and decision analysts swapped ideas and held joint conferences. Each new demonstration of empirical violations of the von Neumann-Morgenstern or Savage axioms, and each new relaxation of the axioms to accommodate those violations, was greeted with intense interest. (Not everyone agreed the axioms needed to be relaxed, but the debates were always exciting.) For rational choice theorists, the attraction of non-expected utility theory was that it promised to yield more flexible and psychologically plausible models of the rational individual that would be “plug-compatible” with earlier models, opening the door to new research in which non-expected-utility-maximizers would be substituted in economic roles formerly occupied by expected-utility-maximizers.

What might be called a “foreshock” of the non-expected utility revolution was the 1977 publication in *JPE* of a paper called “Risk, probabilities, and a new theory of cardinal utility” by Jagdish Handa, which proposed a new theory of utility based on a weakening of the independence axiom. The model in the paper was obviously fatally flawed and provoked a flurry of letters to the editor of the journal—one of which, by Peter Fishburn, was published in 1978—but meanwhile it inspired Fishburn and other researchers (including Mark Machina and John Quiggin, who had also sent letters) to begin searching for better ways to relax the conventional preference axioms.

In 1979, Kahneman and Tversky's “prospect theory” paper appeared in *Econometrica*, presenting a descriptive model of choice under uncertainty that departed from the expected-utility model in several important respects. For one thing, it postulated that the determinant of utility was not final wealth but rather *gains and losses* relative to a reference point. Kahneman and Tversky postulated that subjects were typically risk averse with respect to gains and risk-seeking with respect to losses, a phenomenon they termed the “reflection effect.” (That is, subjects would rather have a certain gain than a uncertain gain with the same expected value, whereas they would rather avoid a certain loss and would instead prefer an uncertain loss with the same expected value.) Second, they postulated that subjects evaluate uncertain prospects using “decision weights” that are distorted versions of the probabilities of outcomes. In particular, they assumed that subjects typically act as if they overestimate very small probabilities and underestimate moderate to large probabilities. Kahneman and Tversky showed that these assumptions were able to explain a variety of well-documented preference anomalies,

¹ Ellsberg freely admitted in his paper that his own experiments were carried out “under absolutely non-experimental conditions... Here we see the advantage of purely hypothetical experiments. In ‘real life,’ you would probably turn out to have a profound color bias which would invalidate the whole first set of trials, and various other biases would show up one by one as the experimentation progressed inconclusively.”

including the Allais paradox. “Prospect theory” became the most widely cited paper ever published in *Econometrica*, an indication of its cross-disciplinary impact—not to mention its accessibility. (The original version of prospect theory was quickly recognized to have a serious flaw: the distortion of outcome probabilities could lead to behavior that violates first-order stochastic dominance, which even psychologists would admit is irrational. This flaw was later corrected in *cumulative prospect theory*, a revised form of the theory in which cumulative rather than non-cumulative probabilities are systematically distorted. This has become the most widely studied non-expected utility model, and it is discussed in more detail below.)

Another significant 1979 publication was the volume *Expected Utility Hypothesis and the Allais Paradox* edited by Maurice Allais and Ole Hagen, which contained a translation of Allais’ original paper illustrating a violation of the independence axiom, together with other papers that raised new questions about the validity of the von Neumann-Morgenstern and Savage axioms. This volume precipitated the founding of a series of bi-annual interdisciplinary meetings on “foundations of utility and risk” at which many of the emerging ideas on non-expected utility were presented and discussed. (Most these meetings were held in Europe, although the fifth meeting in the series was held at Duke in 1990.) Allais’ ideas were once again a focal point of debate in economic theory, and he was awarded the Nobel prize in 1988 (though mainly for his contributions to equilibrium theory, not utility theory).

The idea that preference axioms (especially the independence axiom) ought to be modified so as to have more empirical validity, and the idea that subjects somehow use distorted versions of probabilities in their evaluation of risky prospects, captured the imagination of many researchers and led to the construction of many new representations of preferences. Kahneman and Tversky’s prospect theory was not axiomatically based—it was merely proposed as a descriptive model of not-necessarily-rational choice—and it had some properties that were objectionable from a normative viewpoint. In particular, the type of probability-distortion postulated by Kahneman and Tversky implied that subjects would fail to respect *stochastic dominance*. A risky prospect x is said to “stochastically dominate” another risky prospect y if x yields an equal or greater cumulative probability than y of exceeding any specified target level of wealth. Most researchers consider it an unassailable requirement of rationality that x should be preferred to y if it is stochastically dominant over y . Following the publication of Kahneman and Tversky’s paper, more researchers began to seek representations of preferences that would allow violations of the independence axiom while remaining faithful to other requirements of rationality such as respect for stochastic dominance.

The shot-heard-round-the-world in the non-expected utility revolution was fired by Mark Machina in a lead article published in *Econometrica* in 1982. Machina presented a model of *generalized expected utility* in which preferences are represented by an “expected utility functional” that behaves like expected utility locally but not globally. Machina illustrated his theory using “triangle diagrams” originally introduced by Jacob Marschak, which are now widely known as “Marschak-Machina triangles.” A triangle diagram represents the set of all possible probability distributions over lotteries with three outcomes. By convention, triangle diagrams are drawn so that the vertex in the lower left represents the intermediate outcome; the vertex in the upper left, or “northwest,” represents the best outcome, and the vertex in the lower right, the “southeast,” represents the worst outcome. Under expected utility theory, an

individual's indifference curves with respect to such lotteries must consist of *parallel straight lines* in the triangle diagram, as a consequence of the property of linearity in the probabilities. Under Machina's theory, the indifference curves can have practically any shape—as long as they are “smooth”—although Machina hypothesized that they “fan out” towards the northwest and southeast corners, a pattern that would explain the Allais paradox and a number of other stylized facts.

The following year (1983), Soo Hong Chew published a paper in *Econometrica* that presented an axiomatic model of *weighted utility* in which the independence axiom was replaced by a weaker “betweenness” axiom. Betweenness requires that if $f \succ g$ then $\alpha f + (1-\alpha)g \succ g$ for all α between 0 and 1. That is, if f is preferred to g , then everything “between” f and g should also be preferred to g . (Recall that the usual independence axiom requires that if $f \succ g$ then $\alpha f + (1-\alpha)h \succ \alpha g + (1-\alpha)h$ for an arbitrary third lottery h , which means that if f is preferred to g then everything between f and h is preferred to the same relative interpolation between g and h .) Under Chew's weighted utility theory, indifference curves are *straight lines radiating from a single point* outside and usually southwest of the Marschak-Machina triangle, consistent with the fanning-out hypothesis. (The first version of this model was proposed in a 1979 working paper jointly written with Ken MacCrimmon.)

Machina's and Chew's models both explained the typical preference pattern of the Allais paradox and provided more flexible tools for economic modeling than the standard expected utility model. But they did not capture the idea of distorted probabilities that was one of the intriguing elements of prospect theory. Machina's model also was not derived from preference axioms—rather, it was derived from assumptions imposed on the utility functional that represents preferences—and this eventually came to be perceived as a liability. The idea of distorted probabilities was appealing because it seemed like a natural “dual” explanation of attitudes toward risk. Risky prospects are typically composed of *payoffs* and *probabilities*: if a subject is allowed to subjectively transform the payoff scale by applying a utility function, why shouldn't she also be allowed to transform the probability scale? Intuitively, risk aversion is a response to the presence of *uncertainty*, so it seems plausible that its explanation might somehow involve the decision maker's appreciation of probabilities rather than amounts of money.

Various researchers explored this other avenue, and an axiomatically sound way to represent preferences in terms of distorted probabilities was soon independently discovered by John Quiggin (1982), David Schmeidler (1982/1989), and Menachem Yaari (1987). The key idea in the Quiggin/Schmeidler/Yaari approach is that *cumulative* probabilities, rather than non-cumulative probabilities, are subjected to a systematic distortion. A subject who monotonically distorts the cumulative probability distributions of risky prospects will not violate stochastic dominance. It turns out that the distorted-cumulative-probability representation follows from a relaxation of the Savage axioms in which the independence axiom is replaced by a weaker condition of *comonotonic independence*. Two acts are said to be “comonotonic” if they yield the same ranking of consequences across states. That is, if f and g are comonotonic, and $f(s) \succ f(s')$, then also $g(s) \succ g(s')$, where $f(s)$ denotes the consequence yielded by f in state s . The comonotonic independence axiom states that acts f , g , h , and h' must satisfy the usual independence condition *only if they are comonotonic*. The Quiggin/Schmeidler/Yaari model is variously referred to as “anticipated utility,” “rank-dependent utility,” or “Choquet expected

utility.” (The version of the model that treats decision under uncertainty uses a so-called Choquet integral in the preference representation.) The recent paper by Diecidue and Wakker gives a very lucid explanation of the intuition behind the rank-dependent model and the comonotonic independence condition.

In the early 1990’s, Kahneman and Tversky—with help from Peter Wakker—succeeded in axiomatizing *cumulative prospect theory*, which replaces the distorted probabilities in the original theory with distorted cumulative probabilities, in the manner of rank-dependent utility theory. Cumulative prospect theory still refers to losses and gains relative to a reference point: it is rank *and sign* dependent. Empirical tests of this model in recent years suggest that it may have some descriptive advantages over original prospect theory (Miyamoto and Wakker, *Journal of Risk and Uncertainty* 1997).

Another major branch of non-EU theory consists of theories that relax the *transitivity axiom* rather than the independence axiom. Although transitivity is widely considered to be normatively compelling, there are well-known situations in which many subjects tend to violate it. One pattern of violation occurs in situations where feelings of *regret* play a role in decision making: to a regretful decision maker, what matters is not merely the consequence that is actually obtained, but also the consequence that *might have been obtained* if an alternative decision had been taken. It is easy to construct examples in which a subject will violate transitivity in a sequence of pairwise comparisons because of feelings of potential regret that are induced by way in which alternatives are paired. Models of regretful behavior were proposed independently by Graham Loomes and Robert Sugden (1982, *Economic Journal*) and by David Bell (1982, *Operations Research*). A more general theory of “skew-symmetric bilinear” (SSB) utility, which allows regret and other forms of intransitive preference, was developed by Peter Fishburn in a long series of papers and a 1988 book, *Nonlinear Preference and Utility Theory*.

Another way in which individuals may construct intransitive preferences is by merely counting the number of events or attributes in which one prospect is better than another. A famous example is that of a person who is trying to decide whom to marry among three candidates, on the basis of wealth, beauty, and intelligence. Candidate A ranks 1st on wealth, 2nd on beauty, and 3rd on intelligence. Candidate B ranks 3-1-2, respectively, and candidate C ranks 2-3-1, respectively, on the same attributes. Hence candidate B is better than A on two out of three attributes, candidate C is better than B on two out of three attributes, and candidate A is better than C on two out of three attributes. A person who determines his or her preference between each pair of marriage partners by counting the attributes on which one is better than the other will in this case exhibit intransitivities. (This is actually just the Condorcet paradox of voting cycles under another name. Suppose that “wealth,” “beauty” and “intelligence” are three voters who have different rankings for three candidates for a political office. Then two out of three voters prefer B to A, two out of three voters prefer C to B, and two out of three voters prefer A to C. Hence, pairwise majority voting may lead to cycles rather than a well-defined social preference.)

The same phenomenon can occur when choices are made among risky gambles. A famous (or perhaps infamous) example is the following: suppose that a six-sided die is to be rolled and you can choose among six different gambles whose payoffs depend on the outcome. Gamble 1 yields \$100 if 1 is thrown, \$200 if a 2 is thrown, and so on up to \$600 if a 6 is thrown. Gamble 2 is

similar to gamble 1, except that the payoffs are all “rotated” by one position on the die: you get \$100 if a 2 is thrown, \$200 if a 3 is thrown, etc., and \$600 if a 1 is thrown. Gambles 3 through 6 are defined by successive rotations of the payoffs. Thus, gamble 6 yields \$100 if a 6 is thrown, \$200 if a 1 is thrown, and so on. Consider the following pattern of preferences: gamble 1 is strictly preferred to gamble 2 because it yields a higher payoff in 5 out of 6 cases. Gamble 2 is strictly preferred to gamble 3 by the same reasoning, etc., and finally gamble 6 is strictly preferred to gamble 1, creating a cycle. This example was presented by Fishburn at a national OR/MS conference in 1987, and he confessed that his own personal preferences among these gambles followed exactly this pattern, which is rational according to his own SSB theory. Fishburn’s stark confession caused something of an uproar in the decision analysis community. Despite the fact that SSB utility theory is a mathematically elegant and beautiful generalization of expected utility theory, not many decision analysts would go as far as Fishburn in applying it to their own preferences or those of their clients!

On another front, Fishburn and Irving LaValle (1992, *J. Math. Psych.*) have extensively explored the generalization of expected utility that is obtained by relaxing the *continuity axiom*, sometimes called the “Archimedean” axiom. This leads to a “non-Archimedean decision theory” in which preferences may be lexicographically ordered. For example, suppose that acts lead to consequences with two attributes: a quantity of material wealth and an eternal reward in the life hereafter. In this case, you may prefer (\$ x , heaven) to (\$ y , hell) even if $y \gg x$. But presumably you prefer (\$ x , heaven) to (\$ y , heaven) and you prefer (\$ x , hell) to (\$ y , hell) if $x > y$. Fishburn and LaValle show that such preferences are represented by vector-valued utilities and matrix-valued probabilities. Needless to say, the elicitation of so many parameters is potentially difficult, and applications in areas other than philosophy or theology are somewhat hard to envision. (Even decisions that involve “fates comparable to death” are not necessarily outside the scope of ordinary expected utility theory. It takes a pretty extreme attribute to justify the assumption of lexicographic preferences.)

Just to “complete” the topic of relaxations of the axioms of expected utility theory, we note that the most questionable axiom of the received theory is undoubtedly the assumption of *completeness*, i.e., the assumption that between any two alternatives the decision maker always has at least a weak if not a strict preference. Realistically, the decision maker should be permitted to be occasionally undecided, especially when the set of alternatives is very large and full of unfamiliar—even counterfactual—objects, as is typically the case in formal theories of choice under risk or uncertainty. Dropping the completeness assumption greatly reduces the normative strain on all the other axioms: for example, transitivity and independence are more reasonable requirements when they are applied to strongly-held preferences that the decision maker has asserted voluntarily after careful reflection, as opposed to possibly-unstable preferences elicited through forced answers to artfully framed questions. Robert Aumann published a paper on “Expected utility theory without the completeness axiom” in *Econometrica* in 1962, in which he showed that dropping the completeness axiom (in a von Neumann-Morgenstern framework of objectively known probabilities) leads naturally to a representation of preferences by *convex sets of utility functions* rather than a unique utility function. Similarly, dropping the completeness axiom from subjective probability theory (while ignoring the complications of utility) leads to a representation of beliefs by *convex sets of probability distributions* rather than by a unique distribution. (The literature on sets of subjective probabilities dates back to work by Bernard Koopman in the 1940’s and I.J. Good in the 1950’s,

and this model is now widely used as a foundation for robust methods of Bayesian statistics.) Finally, if you drop the completeness axiom from the theory of subjective expected utility, you get a representation of preferences by *convex sets of state-dependent-utilities*. (The latter representation has been recently explored by Teddy Seidenfeld, Mark Schervish, and Jay Kadane; by David Rios Insua and Jacinto Martin; and by myself.) Models of incomplete preferences (and especially, incomplete beliefs represented by sets of probabilities) began to receive new attention from economists in the mid-1980's, inspired by the work of Truman Bewley on "Knightian" uncertainty. This avenue is now being further explored in a series of papers by Luca Rigotti, Chris Shannon, and others.

Thus, by now, nearly all possible relaxations of the axioms of expected utility and subjective expected utility have been explored. Has this led to the hoped-for new paradigm of choice theory? My impression is that it has not: the old paradigm has been extended, not demolished. You might say that a few new wings have been added to the old house, providing employment for new generations of artisans, but the original structure and its foundation remain in place. Economists have found that many standard results remain valid under more general preference assumptions (as Machina emphasized in his 1982 paper), and the additional parameters of non-expected-utility models allow a somewhat wider range of empirical data to be fitted when necessary. Meanwhile, the phenomenon of nonlinear response to probability has engaged the interest of psychologists, and the study of the "probability weighting function" is a flourishing cottage industry among behavioral decision theorists. But by the end of the 1980's, the proliferation of new utility theories had become, to many, a source of frustration rather than liberation. (Later, we will see that a similar situation developed with respect to refinements of Nash equilibrium in game theory during roughly the same years.) Each of the new axiomatic utility theories explains some subset of the experimental phenomena, but none has proved effective and parsimonious across-the-board. Prospect theory tends to provide the best all-around fit to data from laboratory experiments—which of course it was designed to do—but EU theory does reasonably well in the "interior of the triangle," where probabilities are not too close to zero or one, and it remains on the "efficient frontier" of available theories. (The 1992 paper by Camerer summarizes some of these findings.) Thus far, none of the new theories has succeeded in dislodging subjective expected utility theory as the cornerstone of decision analysis and the economics of uncertainty: most decision analysts and economists drifted back to the "old time religion" (if, indeed, they had ever strayed). In 1989, Ward Edwards organized a conference in Santa Cruz to which leading decision researchers from various camps were invited. He posed two questions:

1. *Do you consider SEU maximization to be the appropriate normative rule for decision making under uncertainty?*
2. *Do you feel the experimental and observational evidence has established the fact that people do not maximize SEU; that is, that SEU maximization is not defensible as a descriptive model of the behavior of unaided decision makers?*

The unanimous response was "yes" to both questions.

Does this mean that all conceivable alternatives to subjective expected utility theory have been explored? I don't think so. The various non-expected-utility theories tweak the von Neumann-

Morgenstern and Savage *axioms* in different directions, but they do not question the deeper foundations of the theory: its focus on *binary preferences* held by *individuals* with respect to abstract *acts* that map states of nature (or which otherwise assign probabilities) to *consequences* whose value is somehow known to be state-independent. But there are grounds for questioning whether binary preferences are psychologically primitive and also for questioning whether individual preferences among abstract acts and consequences are the most useful terms in which to model personal decisions or economic and social institutions. We will return to these questions later

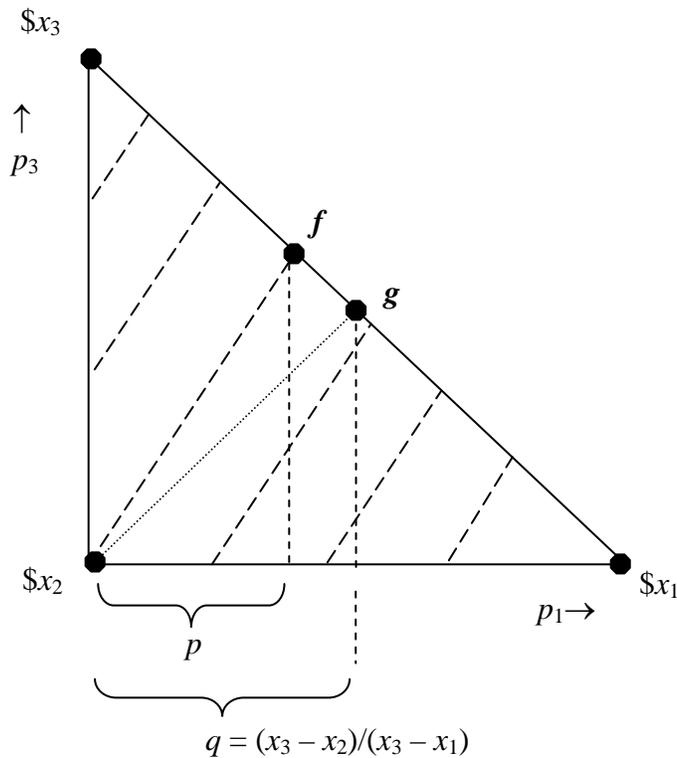
Graphical representations of non-expected-utility preferences under risk

The remainder of this note will discuss the most popular models of non-expected-utility preferences in bit more detail, concentrating on their graphical representations and psychological intuition. In choice theory there has traditionally been a distinction between decision-making under conditions of “risk” (where probabilities are objectively determined) and decision-making under conditions of “uncertainty” (where probabilities must be subjectively determined). This distinction has been preserved in models of non-expected utility preferences. Of course, outside of gambling casinos, game shows, and psychology labs, there is almost no such thing as a decision under risk in real life, but nevertheless much of the literature of choice theory deals with risk rather than uncertainty. Partly this is because risk can be viewed as a simpler, special case of uncertainty (crawling before walking); partly it is because decisions under risk are easier to study in the laboratory (where the experimenter can control the probabilities); and partly it is because some economic theorists are still more comfortable thinking of probabilities in objective terms (i.e., viewing the market as an objective stochastic process).

Models of **non-expected utility preferences under risk** fall into two main categories: (i) those based on the idea of “local” utility functions with indifference curves that “fan out” and (ii) those based on the notion of distorted probabilities, a.k.a. “rank dependence.” The properties of models in category (i) can be illustrated and compared in so-called “triangle diagrams” which depict the decision maker’s indifference curves among probability distributions over a set of 3 elements: a low, medium, and high payoff. In the diagram, the probability distributions are plotted in right-triangular coordinates. By convention, the *lowest* payoff x_1 is associated with the lower-right (southeast) vertex, the *highest* payoff x_3 is associated with the upper-left (northwest) vertex, and the intermediate payoff x_2 is associated with the lower-right (southwest) vertex. The planar coordinates are then p_1 and p_3 , the probabilities assigned to the low/high payoffs x_1 and x_3 , and the remaining probability $1-p_1-p_3$ is implicitly assigned to the intermediate payoff x_2 . In general, if a decision maker’s preferences among probability distributions are complete, transitive, continuous, and monotonic, then they are represented by an ordinal utility function $U(\mathbf{p})$ whose indifference curves can be plotted inside the triangle, as in consumer theory. If she also satisfies the independence axiom, she is an expected utility maximizer and her utility is linear in the probabilities:

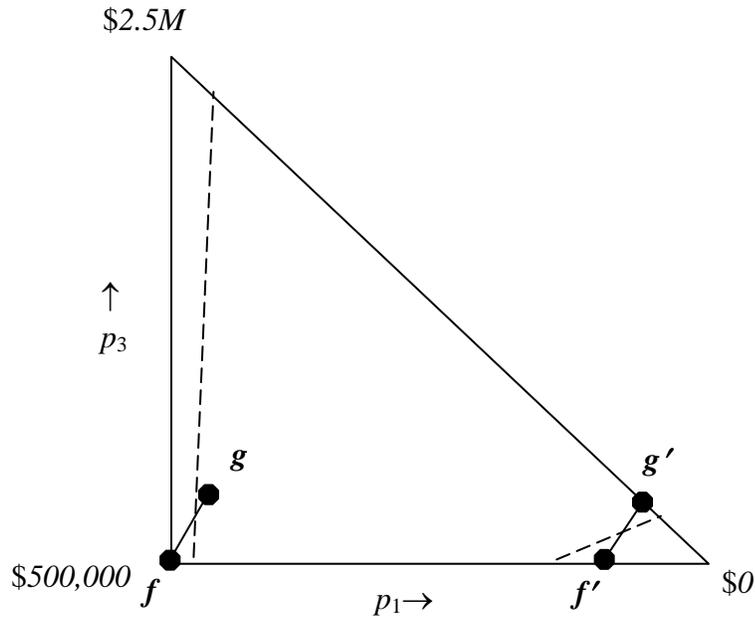
$$U(\mathbf{p}) = p_1u(x_1) + p_2u(x_2) + p_3u(x_3),$$

from which it follows that her indifference curves must be *parallel straight lines* as shown below:

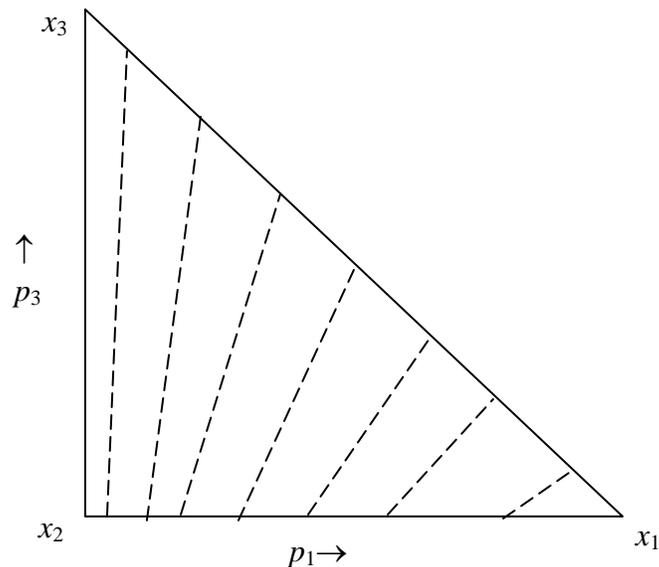


Here, the lottery f assigns probability p to x_1 and $1-p$ to x_3 , and it lies on the same indifference curve as the distribution assigning probability 1 to the intermediate payoff x_2 , hence x_2 is the *certainty equivalent* of f . Without loss of generality, we can assign utility values of 0 and 1 to the low/high payoffs x_1 and x_3 , respectively, and it then follows that $u(x_2) = 1-p$. Once the relative utility of x_2 is known, the utility function is completely specified and the slopes of all the indifference curves are determined. The lottery g , which yields x_1 with probability q and x_3 with probability $1-q$, where $q = (x_3 - x_2)/(x_3 - x_1)$, has an *expected value* of x_2 by construction. Because $p < q$, the decision maker is risk averse—i.e., her certainty equivalent for x_2 is less than its expected value—and her indifference curves are steeper than if she were risk neutral. As p becomes smaller, the individual becomes more risk averse, and her indifference curves become “steeper” in the triangle.

The Allais paradox illustrates a pattern of preferences in which the decision maker’s indifference curves appear to be non-parallel. The lottery f , which yields the intermediate payoff of \$500,000 for sure, is strictly *more* preferred than g , which yields a 0.89 chance of \$500,000, a 0.10 chance of \$2.5M and a 0.01 chance of \$0. Meanwhile, f' , which yields a 0.89 chance of \$0 and a 0.11 chance of \$500,000, is strictly *less* preferred than g' , which yields a 0.90 chance of \$0 and a 0.10 chance of \$2.5M. In a triangle-diagram plot of this situation, the straight lines connecting f with g and f' with g' are parallel, but f evidently lies on a higher indifference curve than g while f' lies on a lower indifference curve than g' , suggesting that the decision maker’s indifference curves look like the dashed lines in the figure:



Various authors, including Machina and Chew, have proposed models of non-expected utility preferences in which indifference curves “fan out” from the southwest corner, consistent with the Allais paradox. Machina assumes that the decision maker has “smooth” (Frechet differentiable) preferences among probability distributions over monetary payoffs, and he shows that this implies that the decision maker must be at least “locally expected utility maximizing” with a utility function for money that may vary (smoothly) with the probability distribution. (Local expected utility maximization arises merely as a byproduct of smoothness, because any smooth function is locally linear.) In the vicinity of probability distribution F , the decision maker acts as if she has a utility function for money $U(x; F)$. This function can be differentiated with respect to its first argument (the amount of money x) to obtain a local Pratt-Arrow measure of risk aversion: $r(x; F) = -U''(x; F)/U'(x; F)$. Machina’s famous “Hypothesis II” states that if F^* stochastically dominates F , then $r(x; F^*) \geq r(x; F)$, i.e., the individual should become more locally risk averse, which implies the pattern of indifference curves noted above: “[I]f the individual satisfies Hypothesis II, stochastically dominating shifts [in probability] represented by upward or leftward movements in the (p_1, p_3) plane will make the local utility function more risk averse and thus raise the slope of the indifference curves, so that the indifference curves will appear ‘fanned out.’”



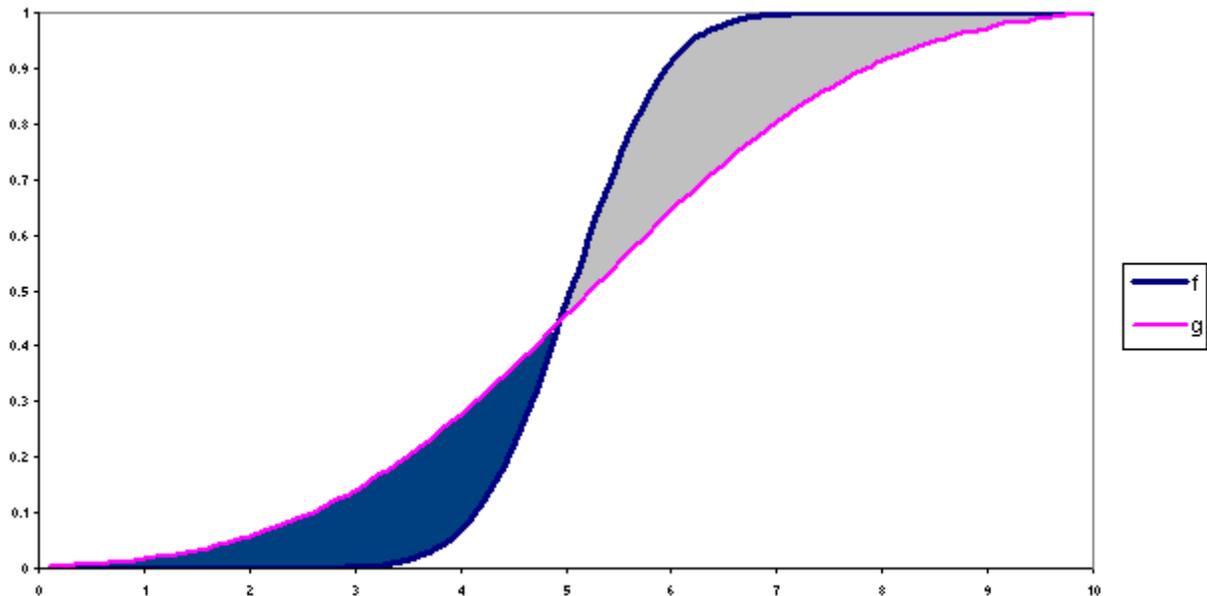
Machina goes on to show that many standard economic results, including various theorems on risk aversion, still hold in this framework and hence are robust to violations of the independence axiom. Chew’s model places more restrictions on preferences, requiring the indifference curves to be straight lines emanating from a point outside the triangle.

More details of these models, as well as empirical evidence for and against the fanning-out hypothesis, are provided in the paper by Camerer from the edited volume based on the 1989 Santa Cruz conference organized by Ward Edwards. It gives a good summary of various non-expected-utility models that had been developed by the end of the 1980’s, as well as a review of their empirical successes and failures in behavioral experiments. It includes some very nice examples of triangle diagrams illustrating the theoretical properties of different non-expected utility models, and it then shows how the contours of the triangle have been “explored” by various experimental investigations.

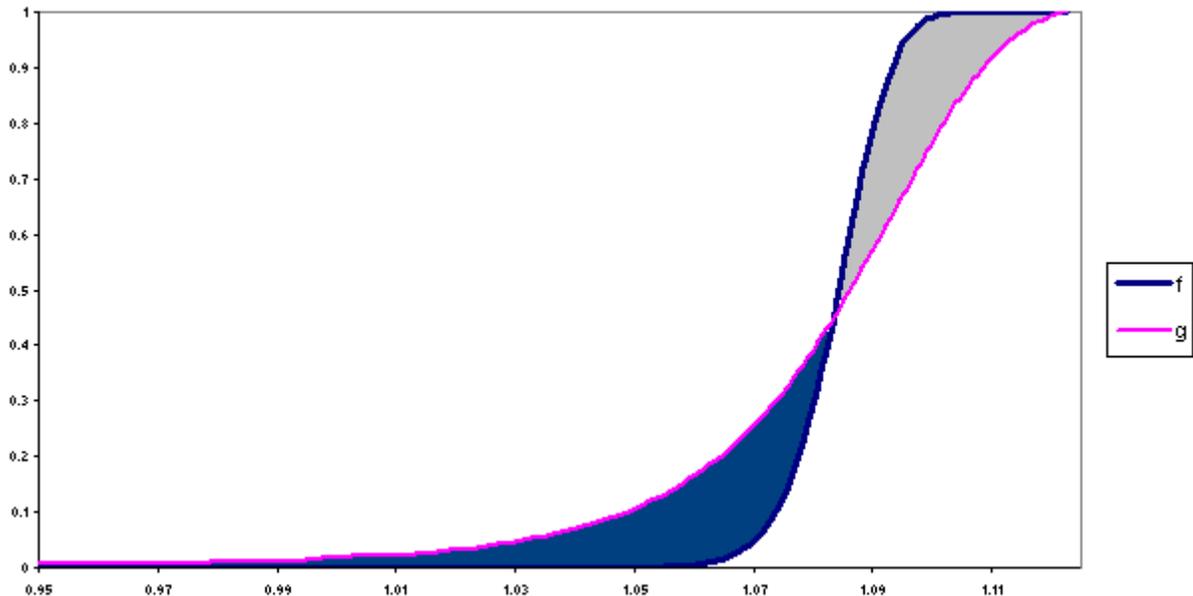
You may notice that the fanning-out hypothesis makes more sense along the southern border of the triangle than in the far north. If the decision maker is in the southeast corner, where nearly 100% of the probability is on the worst possible outcome, she has almost nothing and therefore has almost nothing to lose by gambling. As her distribution moves west toward the southwest corner—i.e., from nearly 100% x_1 to nearly 100% x_2 , she acquires an intermediate payoff with near-certainty, and it is reasonable to imagine that she might become more conservative in her gambling behavior. But at the far northern tip of the triangle she is practically assured of the best outcome, and it is not clear why she would not start to become risk neutral. For this reason, other authors have suggested that the indifference curves should “fan in” at the north and/or bend to the right along the hypotenuse. The **rank-dependent model** has the latter property

The intuition behind the rank-dependent model is most easily illustrated in graphs of *cumulative probability distributions* for decisions with continuously-distributed monetary payoffs. The

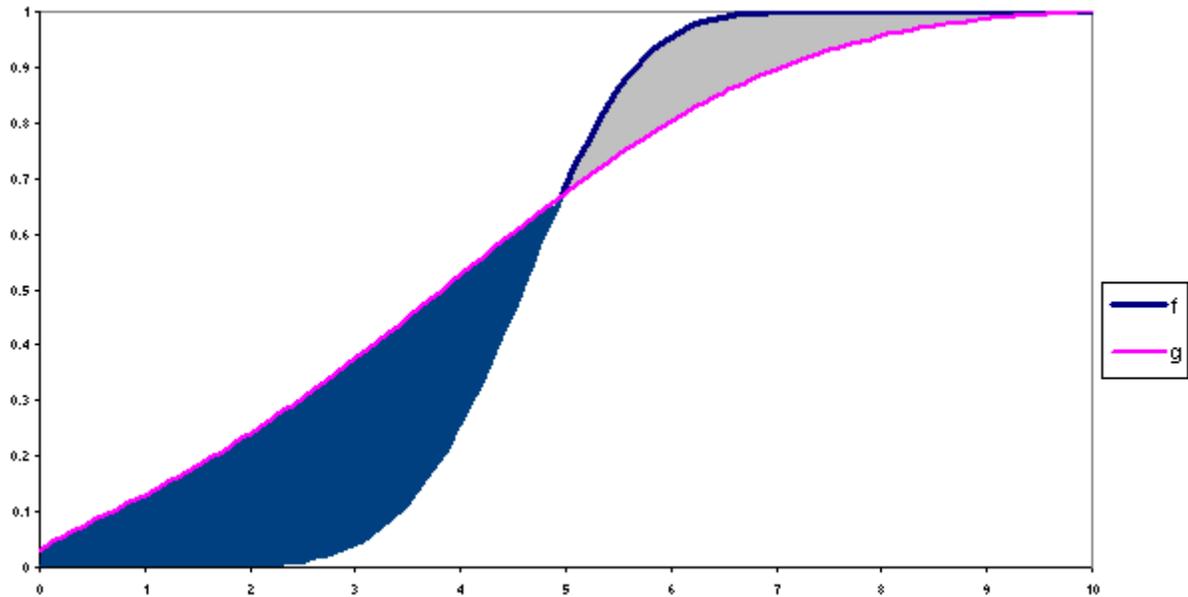
graph below shows a comparison of the cumulative distributions of two decisions f and g , where f 's payoff has a lower variance but g 's payoff has a slightly higher mean. **If the distributions are bounded below at zero, the expected payoff of each decision is the area to the left of its cumulative distribution curve and to the right of zero (an under-appreciated mathematical fact!).** Therefore, we can determine which is the expected-value maximizing decision by comparing the two shaded regions in the graph below: the dark-shaded region represents area to the left of f that is not to the left of g , and the light-shaded region represents area to the left of g that is not to the left of f . As you can see, the light-shaded region is a little bigger, so g indeed has a higher expected payoff.



Now suppose that the decision maker is a risk-averse expected utility maximizer. Then she will not wish to compare the expected *values* of the payoffs, but rather their expected *utilities*. This can be illustrated by redrawing the graph so that the x -axis is scaled in utility units instead of money. The utility scale is nonlinearly distorted relative to the monetary scale so that it magnifies the *left side* of the graph. The result is shown in the picture below. As you can see, the curves have been stretched out toward the left, magnifying what happens at low payoff values. We can now compare expected *utilities* by comparing the relative sizes of the two shaded regions. As you can see, the darker shaded region is now bigger, indicating that the more “conservative” decision f has a higher expected utility than the more “risky” decision g .



Thus, risk aversion is conventionally modeled by a *subjective nonlinear distortion of the payoff scale* on a plot of the cumulative distributions. Over the years, many theorists have criticized the psychological validity of this representation of risk aversion, in which the decision maker avoids risk only because she has diminishing marginal utility for money. It is commonly observed that both individuals and firms sometimes act in a risk averse manner with respect to gambles whose payoffs are small enough that nonlinear utility for money should not have any effect. What if the decision maker's utility for money is really *linear*, but she just doesn't like *risk*? The presence of risk is signified by the fact that some payoffs have probabilities other than 0 or 1. Perhaps the decision maker's risk attitude is in some way explained by her subjective appreciation of the probabilities. In particular, suppose that she subjectively distorts the *probability* scale instead of (or perhaps in addition to) the monetary scale? A nonlinear distortion of the probability scale can be shown on the same kind of graph. The following figure shows the graph of the cumulative distributions of f and g with the probability scale nonlinearly transformed so that the probabilities of low payoffs are amplified and probabilities of high payoffs are attenuated. This stretches the *lower half* of the graph (whereas previously we stretched the *left half*).



As you can see, the dark-shaded region is again larger than the light-shaded region, indicating that the more conservative decision f is preferred despite its lower expected value—and despite the fact that the decision maker now is assumed to have linear utility for money.

Yaari’s 1987 paper on the “dual” theory of choice under risk gives a very elegant derivation of this model in which the independence axiom is applied to payoff mixtures instead of probability mixtures, with the result that the probability scale ends up being distorted instead of the monetary scale. Recall that we previously derived de Finetti’s theorem on subjective probabilities by “dually” applying the independence axiom to payoff mixtures rather than probability mixtures. This is essentially what Yaari does. He also adds the assumption that the states have objective probabilities and that preferences respect stochastic dominance. It follows that the subjective probabilities that the decision maker ultimately uses to evaluate gambles must be monotonic transformations of the objective ones. The psychological intuition behind this model—namely that the decision maker pays more “attention” to events with low-ranked payoffs—is also very nicely explained in the paper by Diecedue and Wakker.

More general models of non-expected utility preferences under uncertainty

Although models of non-expected utility preferences under risk are mathematically elegant and easily lend themselves to testing in laboratory experiments involving objective probabilities, they are ultimately less interesting than models of **non-expected utility preferences under uncertainty**, since they are inapplicable to “real” events and they are unable to explain phenomena such as Ellsberg’s 2-color paradox, in which individuals are averse to betting on events whose probabilities are “ambiguous.” There are four main types of non-expected utility models under uncertainty: (i) Schmeidler’s Choquet expected utility (and its close cousin, Gilboa & Schmeidler’s maxmin expected utility), which is a generalized form of the rank-dependent model; (ii) cumulative prospect theory, which generalizes Choquet expected utility to include reference-point effects and sign-dependent as well as rank-dependent probabilities; (iii) “probabilistically sophisticated” non-expected utility preferences, which is Machina and

Schmeidler's generalization of Machina's model of choice under risk; and (iv) models based on state-preference theory, which are historically the oldest models but have attracted new interest lately.

The **Choquet expected utility** model (also nicely discussed in the paper by Diecedue and Wakker) is a generalization of Savage's subjective expected utility model in which the independence axiom is replaced by the weaker condition of *comonotonic independence*, which applies only to pairs of acts which yield the same ranking of states in terms of the desirability of the consequences to which they lead. More precisely, f and g are defined to be comonotonic acts if the state that yields the most-preferred consequence under f also yields the most-preferred consequence under g , and similarly for the state that yields the second-most-preferred consequence, and so on down the line. The important feature of comonotonic acts is that, since their payoffs have a rank-correlation of 100%, they cannot be used to hedge away each other's risks through the formation of compound acts. The comonotonic independence axiom states that three acts f , g , and h must satisfy the independence axiom *only* if they are comonotonic. Recall that the independence axiom requires

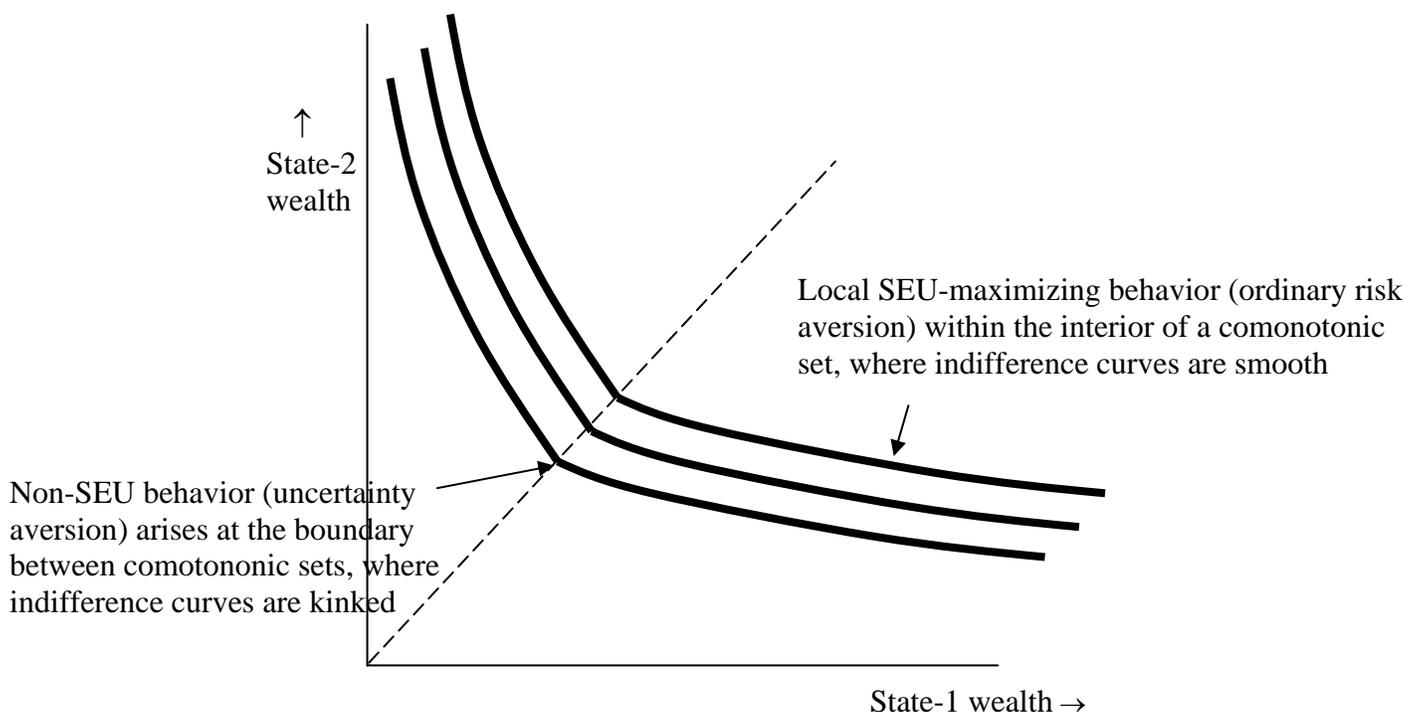
$$f \succcurlyeq g \Leftrightarrow Af + (1-A)h \succcurlyeq Ag + (1-A)h$$

for any non-null event A . If h is comonotonic with f and g , this axiom is more psychologically plausible, because mixing f and g with h cannot not hedge away either of their respective risks. (Recall that the essence of the Allais paradox was that, for some choices of h , the compound act on one side of this relation could be turned into a "sure thing" while the other side remained risky.)

Preferences that obey Schmeidler's axioms are represented by a unique utility function over consequences combined with a *non-additive probability distribution* known as a "Choquet capacity." A capacity is a function W that merely assigns greater weight to event A than event B when B is a subset of A , while a probability measure must assign weights that are additive across mutually exclusive events. Schmeidler referred to this model as "subjective probability and expected utility without additivity," although it has become more commonly known as *Choquet expected utility* (CEU).

Replacing the standard independence axiom with the comonotonic independence axiom has an effect that is most easily visualized in the multidimensional payoff space of state-preference theory. In payoff space, the sets of comonotonic acts are convex cones that meet along the "45 degree certainty line" or its higher-dimensional generalization. For example, in 2-dimensional space, the sets of comonotonic acts are two half spaces that meet along the line $y=x$. In three-dimensional space they are six wedges that meet along the line $z=y=x$, and so on. Within each comonotonic set, the decision maker must behave exactly like an expected utility maximizer, because she obeys all the Savage axioms locally. At the boundary between two comonotonic sets, her indifference curves are *kinked* if the events in question are regarded as "uncertain" rather than merely "risky", as illustrated in the figure below, because her apparent subjective probabilities change discontinuously when that boundary is crossed. When the decision maker is straddling such a kink, as in the Allais and Ellsberg experiments, she exhibits "first-order risk aversion" toward the uncertain events, which means that she is risk averse even for very small bets. By comparison, the usual sort of risk aversion that is quantified by the Pratt-Arrow

measure is “second-order risk aversion,” which is exhibited only when bets are large enough for the second derivative of the utility function to exert an effect. The CEU decision maker is second-order risk averse only with respect to acts that lie within the interior of the same comonotonic set, where her indifference curves are smoothly curved in the usual way.



A close relative of the Choquet expected utility model is the **maxmin expected utility** (MEU) model developed by Gilboa and Schmeidler (*J. Math. Econ.* 1989). In the MEU model, the decision maker’s preferences are represented by a unique von Neumann-Morgenstern utility function u and a convex set C of subjective probability distributions. An act f is preferred to another act g if and only if the minimum expected utility over all probability distributions in the set C at least as great for f as for g . Thus, it is as if the decision maker is unsure about which subjective probability distribution to use (similar to the subjective probability model with incomplete preferences, in which beliefs are also represented by a convex set of distributions), and she has a “pessimistic” (i.e., uncertainty averse) attitude which leads her to value every act at its minimum possible expected utility over the set C of possible probability distributions. The relation between CEU and MEU is as follows: if the Choquet capacity is *convex* (that is, if $W(A \cup B) \leq W(A) + W(B)$ for all events A and B), then the CEU model is equivalent to an MEU model in which the set C of probability distributions is the set of all distributions p such that $W(A) \leq p(A)$ for every event A , where W denotes the Choquet capacity (i.e., the non-additive probability measure). However, the MEU model is not completely “nested” inside the CEU model: depending of the shape of the set C , it is possible that an MEU model does not correspond exactly to a CEU model. Axiomatically, the MEU model (like the CEU model) is derived from a weakening of the Anscombe-Aumann axioms of horse lotteries, which were discussed in class 2 as a simpler way to derive the SEU model. Under MEU, the independence

axiom (A2 in the notes from class 2) is weakened so that it applies only to *constant* horse lotteries, i.e., horse lotteries that yield the same objective probability distribution over prizes in every state of the world. In addition, an axiom of *uncertainty aversion* is explicitly included. The latter axiom states that if $f \sim g$, then $\alpha f + (1-\alpha)g \succcurlyeq f$, in other words, the decision maker prefers to hedge her uncertainty by forming objective mixtures between otherwise-equivalent acts. (This is essentially the same definition of uncertainty aversion used by Schmeidler in his 1989 CEU paper, although there it was merely a definition, not an axiom.) The MEU model is somewhat easier to visualize than the CEU model: uncertainty is represented by a *set* of probability distributions rather than a unique distribution, and aversion to such uncertainty is represented by the pessimistic operation of valuing an act at its minimum possible expected utility when its expected utility is not uniquely determined. (By comparison, the CEU model substitutes the somewhat tricky concept of a “capacity” for the better-understood concept of a probability distribution.)

Machina and Schmeidler’s model of “**probabilistically sophisticated**” **non-expected-utility preferences** is an extension of Machina’s earlier work to the setting of uncertainty. Machina’s original model applied to choices under risk, but any model of choice under risk can be extended to choices under uncertainty if the decision maker can be assumed to have subjective probabilities that behave just like objective probabilities. This is what Machina and Schmeidler do: they start from Savage’s model and drop the independence axiom to allow for non-expected-utility preferences, but at the same time they strengthen the comparative probability axiom (Savage’s P4) so that it will still yield a unique ordering of events by probability even without independence. Machina and Schmeidler replace Savage’s P4 with a “strong” comparative probability axiom, which they call P4*, which has a little of the flavor of the independence axiom thrown in.

Schmeidler’s Choquet model, Gilboa-Schmeidler’s maxmin model, and Machina-Schmeidler’s probabilistically sophisticated model are all capable of explaining the phenomenon of uncertainty aversion revealed by Ellsberg’s 2-color and 3-color paradoxes. The CEU and MEU models do so by assigning non-additive probabilities (or indeterminate probabilities) to events that depend on the colors of balls whose proportions are unknown. This explanation of the paradoxes requires the decision maker to start from a position of constant (i.e., riskless) wealth, so that she is straddling the kink in the indifference curve that occurs at the boundary between two comonotonic sets. If she is somewhere else in payoff space, well away from the kink, she will act like a local subjective-expected-utility maximizer with additive probabilities and state-independent utility and won’t exhibit paradoxical behavior. This will be true if she has significant prior stakes in the event in question. Thus, the CEU/MEU explanation of the paradoxes is “special” rather than general—it only works under rather restrictive initial conditions, namely when the decision maker has no prior stakes in events. (However, the restrictions are admittedly very convenient for empirical testing in the laboratory, where the events are artificial and hence of no prior interest to the decision maker. The Allais paradox involves high-stakes gambles that would presumably straddle two different comonotonic sets regardless of initial wealth, although there are low-stakes versions of it that would not.)

The Machina-Schmeidler model of probabilistically sophisticated preferences accommodates uncertainty aversion in a more indirect manner, as discussed in a recent paper by Larry Epstein (1999, *Review of Economic Studies*). In the spirit of an earlier paper by Yaari (1969), Epstein

first gives definitions of “comparative” risk and uncertainty aversion between individuals, which he then uses to state definitions of “absolute” risk aversion and uncertainty aversion. Decision maker 1 is defined to be *more risk averse* than decision maker 2 if, whenever 1 prefers a *constant* act x to a general act f , then 2 also prefers it. (In other words, 2 never takes a risk that 1 does not take.) A decision maker is then defined to be *risk averse* if there is some risk neutral decision maker than whom she is more risk averse. Epstein next assumes that there are two kinds of acts: ambiguous and unambiguous. Decision maker 2 is defined to be *more uncertainty averse* than decision maker 1 if, whenever 1 prefers an unambiguous act h to a general act f , then 2 also prefers it. (In other words, 2 never chooses ambiguity when 1 does not.) A decision maker is then defined to be *uncertainty averse* if there is some probabilistically sophisticated decision maker than whom she is more uncertainty averse. Note that the definition of risk aversion depends critically on prior notions of risk neutral behavior and riskless acts, while the definition of uncertainty aversion depends critically on prior notions of probabilistic sophistication and unambiguous acts. Savage’s “constant” consequences play a key role here.

Cumulative prospect theory: rank and sign dependence

Kahneman and Tversky’s prospect theory is based on three key ideas. First, the decision maker’s preferences are defined with respect to “prospects”, which are state-dependent changes in wealth relative to the current status quo, rather than “acts”, which are state-dependent final wealth positions. Second, the decision maker may have different risk attitudes with respect to gains and losses. Third, the decision maker responds to uncertainty by assigning to each state a “decision weight” which is a distorted version of the state’s objective or subjective probability. When comparing prospects that rank the states in the *same order* (with respect to the desirability of outcomes they lead to) and which also have the *same pattern of signs* (gains or losses) across states, the decision maker uses an additively separable utility function:

$$U(\mathbf{x}) = \pi_1 v(x_1) + \dots + \pi_n v(x_n)$$

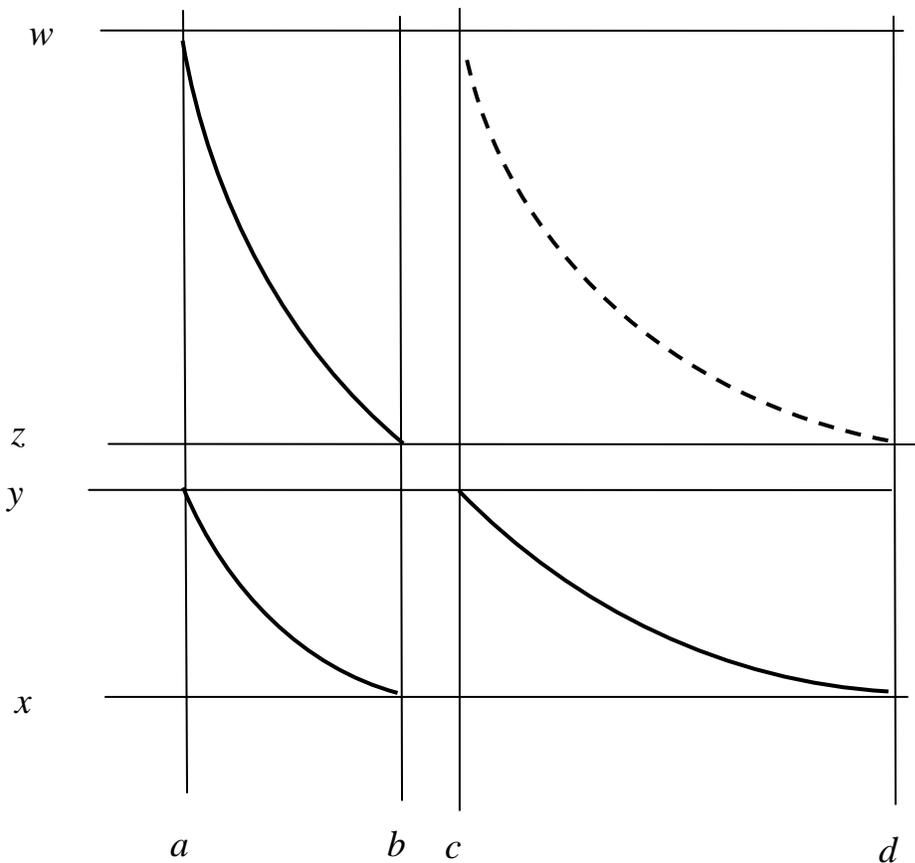
where v is a value function that is typically concave (risk averse) for gains, convex (risk seeking) for losses, and also kinked downward at the origin so that the decision maker exhibits first-order risk aversion for very small gambles (i.e., she strictly prefers a sure thing to any fair gamble). However, the decision weights (π_1, \dots, π_n) are generally both rank and sign dependent, i.e., they may vary according to the pattern of signs in the vector \mathbf{x} as well as the ordering of the states it induces. (Actually, there is also a fourth key idea, namely that the evaluation of prospects is preceded by an “editing” phase in which the decision maker determines the relevant status quo and simplifies the descriptions of the prospects if necessary. The editing phase potentially explains other behavioral phenomena such as “framing effects,” but is outside the scope of the current discussion.)

The paper by Wakker and Tversky (*JRU* 1993) presents an elegant axiomatization of cumulative prospect theory that makes clear its relation to subjective expected utility theory and Choquet expected utility theory. In fact, the paper axiomatizes all the theories at once, as a kind of unified theory of choice under uncertainty. The key tool is the axiom of *generalized triple cancellation*, which (as we saw in class 1) is the essential property of preferences that leads to an additive cardinal utility representation. Here is the GTC axiom once again. In the present setting, the acts \mathbf{x} , \mathbf{y} , etc., are vectors representing amounts of money to be received in different

states of the world (rather than amounts of different commodities to be received simultaneously). Let $b x_i$ now denote the modified act that yields the amount of money b in state i and otherwise yields the same amounts of money in all states as act x .

Generalized Triple Cancellation: $b x_i \leq a y_i$ and $d x_i \geq c y_i$ and $b z_i \geq a w_i \Rightarrow d z_i \geq c w_i$

In words, if trading b for d is as at least as good as trading a for c in state i when the starting point is a comparison of $b x_i$ against $a y_i$, then trading b for d cannot be strictly worse than trading a for c when the starting point is a comparison of $b z_i$ against $a w_i$. Here (again) is a picture that illustrates the special case of GTC that obtains with two coordinates when all of the relations are indifference: the dotted indifference is implied by the three solid ones. This is precisely the property that indifference curves must have in order to possess an additive representation.



The GTC axiom embodies a state-specific form of **tradeoff consistency**: if the tradeoff between b and d is “greater than” the tradeoff between a and c in state i against the background of one comparison of acts, then the same relation between tradeoffs must hold (in state i) against the background of any other comparison of acts. Together with the usual ordering axioms of completeness, transitivity, reflexivity, and continuity, the GTC axiom implies that preferences are represented by an additively separable cardinal utility function:

$$U(\mathbf{x}) = u_1(x_1) + \dots + u_n(x_n).$$

This representation is **state-dependent expected utility without unique determination of probabilities**, since probabilities of states cannot be uniquely factored out of the possibly-state-dependent utility functions for money $\{u_i\}$. To get to the standard SEU model, we need to strengthen the GTC axiom so that the utility functions for money are proportional in all states. Following Wakker (1989) and Wakker and Tversky (1993), we define a stronger “cross-state” notion of tradeoff consistency:

Tradeoff Consistency: $b x_{\cdot i} \leq a y_{\cdot i}$ and $d x_{\cdot i} \geq c y_{\cdot i}$ and $b z_{\cdot j} \geq a w_{\cdot j} \Rightarrow d z_{\cdot j} \geq c w_{\cdot j}$

(Look closely and note that the subscript j rather than i appears in the last four terms.) Thus, TC requires that if the tradeoff between b and d is “greater than” the tradeoff between a and c in state i against the background of one comparison of acts, then the same relation between tradeoffs must hold *in every state* against the background of any other comparison of acts. The TC axiom together with the usual ordering axioms implies that preferences are represented by a utility function of the SEU form:

$$U(\mathbf{x}) = \pi_1 v(x_1) + \dots + \pi_n v(x_n).$$

where (π_1, \dots, π_n) is a unique vector of “decision weights” that behave exactly like subjective probabilities (i.e., they are non-negative and sum to 1) and v is a state-independent utility function for money that is unique up to positive affine transformations. Thus, at one stroke, the TC axiom accomplishes the same effects as Savage’s P2-P3-P4, albeit in a less general setting of consequences that are monetary in nature (or at least elements of a connected topological space). Strictly speaking, we still cannot say that the decision weights are the decision maker’s “true” subjective probabilities because we cannot rule out the possibility that the “true” utility functions for money have state-dependent scale factors.

Now, the Choquet expected utility model (a.k.a. “cumulative utility”) and the cumulative prospect theory model can be obtained by merely restricting the scope of the tradeoff consistency axiom. In the case of Choquet expected utility, the axiom is restricted to comparisons of acts that are *comonotonic*, i.e., which rank the states of nature in the same order in terms of the magnitudes of the payoffs to which they lead. That is, two acts \mathbf{x} and \mathbf{y} are comonotonic if $x_i > x_j$ wherever $y_i > y_j$ for any two states i and j .

Comonotonic Tradeoff Consistency: $b x_{\cdot i} \leq a y_{\cdot i}$ and $d x_{\cdot i} \geq c y_{\cdot i}$ and $b z_{\cdot j} \geq a w_{\cdot j} \Rightarrow d z_{\cdot j} \geq c w_{\cdot j}$ whenever $\{b x_{\cdot i}, a y_{\cdot i}, d x_{\cdot i}, c y_{\cdot i}\}$ are comonotonic and $\{b z_{\cdot j}, a w_{\cdot j}, d z_{\cdot j}, c w_{\cdot j}\}$ are comonotonic.

The CTC axiom together with the usual ordering axioms implies that preferences are represented by a utility function of the CEU form:

$$U(\mathbf{x}) = \pi_1(\mathbf{x})v(x_1) + \dots + \pi_n(\mathbf{x})v(x_n),$$

in which the decision weights $\{\pi_i(\mathbf{x})\}$ depend to some extent on the act \mathbf{x} . In particular, the decision weights depend on the *comonotonic set* in which the act lies—i.e., all acts in the same

comonotonic set are evaluated with the same set of decision weights—and furthermore the decision weights are derived from a Choquet capacity, as discussed above.

In the case of cumulative prospect theory, the objects of choice are “prospects” (distributions for changes in wealth) rather than “acts” (distributions of final wealth), so the zero point and the signs of the payoffs are non-arbitrary. The scope of the tradeoff consistency axiom is further restricted to comparisons of acts that are *sign-comonotonic*, i.e., not only comonotonic but also having identical patterns of positive and negative signs. That is, two acts \mathbf{x} and \mathbf{y} are sign-comonotonic if $x_i > x_j$ wherever $y_i > y_j$ for any two states i and j , and also x_i and y_i have the same sign for every i . In these terms, we have:

Sign-Comonotonic Tradeoff Consistency: $b x_{-i} \leq a y_{-i}$ and $d x_{-i} \geq c y_{-i}$ and $b z_{-j} \geq a w_{-j} \Rightarrow d z_{-j} \geq c w_{-j}$ whenever $\{b x_{-i}, a y_{-i}, d x_{-i}, c y_{-i}\}$ are sign-comonotonic and $\{b z_{-j}, a w_{-j}, d z_{-j}, c w_{-j}\}$ are sign-comonotonic.

The SCTC axiom together with the usual ordering axioms (and one additional technical condition called “gain-loss consistency”) implies that preferences are represented by a utility function of the CPT form:

$$U(\mathbf{x}) = \pi_1(\mathbf{x})v(x_1) + \dots + \pi_n(\mathbf{x})v(x_n),$$

in which the decision weights $\{\pi_i(\mathbf{x})\}$ now depend on the *sign-comonotonic set* in which the act \mathbf{x} lies, and they are not only derived from a Choquet capacity, but decision weights for positive and negative payoffs are derived from *two* separate Choquet capacities, one for positive payoffs and one for negative payoffs. (Whew.) The function v is unique up to a positive scale factor (i.e., it is measurable on a *ratio scale*, not an interval scale) and is usually called a “value function” rather than a utility function. Geometrically, the CPT model yields indifference curves that are kinked along the 45-degree certainty line and also kinked at the boundaries between quadrants, which are boundaries of sign-comonotonic sets. Within a given sign-comonotonic set, the decision maker behaves just like an SEU-maximizer. Violations of independence are exhibited only in comparisons among prospects that lie in different sign-comonotonic sets.

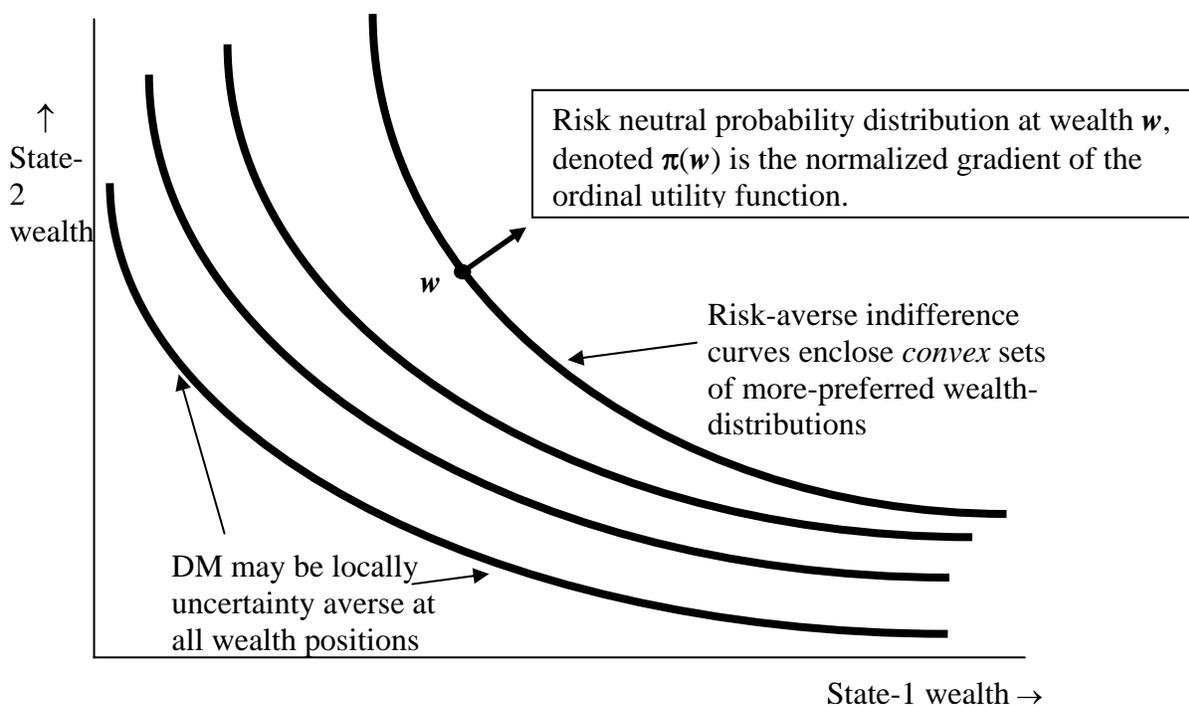
The cumulative prospect theory model, in its full generality, has a large number of free parameters. In practice, it is usually applied to laboratory experiments in which the payoffs have objective probabilities, and the decision weights are assumed to arise from distortions of the cumulative probability distribution, with a simple functional form assumed for the distortion. The value function, as well, is assumed to have a simple functional form that exhibits risk aversion for gains, risk seeking for losses, and a kink at the origin to produce loss aversion (i.e., first-order risk aversion for very small gambles). A handy calculator for CPT, showing the functional forms assumed by Tversky and Kahneman, can be found at this address: http://psych.fullerton.edu/mbirnbaum/calculators/cpt_calculator.htm.

State-preference theory revisited: “smooth” non-EU preferences

The preceding three models (CEU, probabilistically sophisticated choice, and CPT) have many nice features, but they also share an important liability: they do not address the problems in Savage’s model that are even more profound than empirical violations of the independence

axiom, namely the problem of separating probabilities from utilities (embodied in Savage’s axioms P3 and P4) and the problem of constructing a “small world” in which there are “constant” consequences. These models struggle heroically to maintain a clean separation between beliefs and values even as the belief-model and the value-model get messier. The papers by Aumann and Shafer point out the difficulty of identifying consequences whose utility is the same in every state of the world and the difficulty of determining probabilities uniquely from preference data. These tasks do not get any easier when fewer restrictions are placed on preferences. To deal with these problems, we can turn once again to **state-preference theory**, which is an application of plain old 19th-Century ordinal utility theory to preferences under uncertainty, incorporating the Arrow-Debreu trick of defining commodities to be state-contingent.

Recall that in state-preference theory we merely assume that the decision maker has preferences over state-contingent wealth that are complete, transitive, and “smooth” in an appropriate sense. We can visualize the decision maker’s preferences in terms of smooth indifference curves in payoff space, whose normalized gradient at any point is the decision maker’s local risk neutral probability distribution.



If the decision maker also satisfies the independence axiom (or generalized triple cancellation), her ordinal utility function must be additively separable across mutually exclusive events, which places some restrictions on her risk neutral probabilities. For example, in two dimensions the risk neutral probabilities must satisfy the so-called “rectangle property.” If a rectangle is drawn in the (w_1, w_2) plane, then the product of the risk neutral odds ratio π_1/π_2 (which is the marginal rate of substitution between wealth in state 1 and wealth in state 2) between one pair of opposing

corners of the rectangle must be the same as the product of the odds ratio between the other pair of opposing corners. In higher dimensions, the conditional risk probabilities that would be obtained given the occurrence of an event E must be independent of the distribution of wealth in event not- E . (In other words, the marginal rate of substitution of wealth between states i and j should depend only on the current wealth in those two states, not on wealth in other states.) If the decision maker's preferences not only satisfy the independence axiom but also satisfy Savage's axioms P3 and P4 (or tradeoff consistency), which guarantee that they can be represented by a unique probability distribution and state-independent utility function, then her risk neutral probabilities must equal her subjective probabilities at all wealth distributions which are constant (i.e., which lie along the 45-degree certainty line). But in general, the decision maker does not need to satisfy these additional axioms: her indifference curves can have practically any (convex) shape and her risk neutral probabilities can vary with wealth in a more-or-less arbitrary manner.

In the last few years, a number of authors (including myself) have proposed very similar models of “**smooth ambiguity-averse preferences**” in which ambiguity can be represented by a second-order probability distribution together with a **second-order utility function** that represents aversion to the second-order uncertainty. (Klibanoff, Marinacci, and Mukerji 2005, Nau 2001 & 2006, Chew and Sagi 2003; and by Ergin and Gul 2004). In these models it is as if the decision maker evaluates an act by a two-step process. First, she envisions a set of possible first-order probability distributions for states of the world, and for each first-order distribution \mathbf{p} , she computes the expected utility of an act using a first-order utility function u . Then she uses a second-order utility function v to perform some possibly-nonlinear transformation of each first-order expected utility, and finally she uses the second-order probability distribution \mathbf{q} (which assigns probabilities to each of the first-order distributions) to compute the expected value of the second-order utility. If the second-order utility function is concave, it will penalize uncertainty about first-order probabilities in the same way that a concave first-order utility function penalizes uncertainty about states of the world. KMM's version of the model is set in a framework similar to that of Anscombe-Aumann, which includes objectively randomized acts, while mine is set in the state-preference framework and does not necessarily separate probabilities from utilities. The key thing is that the qualitative properties of this model do not depend on knowing the location of a 45-degree certainty line. Rather, they depend only on how the indifference curves curve in different directions.

For simplicity, suppose that the set of states of nature and the set of possible first-order probability distributions are both discrete. Suppose there are N states of the world and M first-order probability distributions. Let \mathbf{p}_m denote the m^{th} first-order distribution, which assigns probability p_{mn} to state n , let \mathbf{q} denote the second-order distribution, which assigns probability q_m to the first-order distribution \mathbf{p}_m , and let the first- and second-order utility functions be denoted by u and v respectively. These functions do not necessarily need to be state-independent, but for simplicity I will assume that they are here. Then the utility of an act \mathbf{x} is computed as follows:

$$U(\mathbf{x}) = \sum_{m=1}^M q_m v \left(\sum_{n=1}^N p_{mn} u(x_n) \right)$$

It is easy to use this model to explain Ellsberg's paradox. Suppose that u is linear, i.e., the decision maker is risk neutral when comparing gambles that are evaluated with a known

probability distribution, particularly when the payoffs are as small as hundreds of dollars. Now suppose that v is concave. Such a decision maker doesn't mind taking risks *per se*, but she hates not knowing the correct probabilities. To explain Ellsberg's 2-urn paradox, suppose that the subject thinks that there are either 100 red balls or 100 black balls in the second urn with equal probability (i.e., there is no middle ground). This means that her possible first-order probability distributions over the four states are $\mathbf{p}_1 = (\frac{1}{2}, 0, \frac{1}{2}, 0)$ and $\mathbf{p}_2 = (0, \frac{1}{2}, 0, \frac{1}{2})$, and that she assigns equal second-order probabilities to these first-order distributions, i.e., $\mathbf{q} = (\frac{1}{2}, \frac{1}{2})$. Then f and f' have first-order expected payoffs of \$50 under both \mathbf{p}_1 and \mathbf{p}_2 so their second-order expected utility is $v(\$50)$. Whereas, the first-order expected payoffs of g are \$100, and \$0, which occur with equal probability, and vice versa for g' . The second-order expected utilities of the latter two acts are therefore both equal to $\frac{1}{2}v(\$0) + \frac{1}{2}v(\$100)$, which is strictly less than $v(\$50)$ if v is strictly concave, which rationalizes a preference for betting on the "known" urn rather than the "unknown" urn regardless of the winning color.

Note that it is the presence of a second-order *utility* function that makes the second-order *probabilities* behaviorally meaningful. It is always possible to construct hierarchies in which higher-order probabilities are used to model uncertainty about the parameter values of lower-order probabilities for events that can be observed. However, at the end of the day, the higher-order uncertainty can be integrated out, leaving only first-order probabilities attached to the observable events. But when a nonlinear second-order utility function is inserted between the first- and second-order probabilities, the hierarchy can no longer be collapsed in this way.

The second-order probability and utility model is very general, but the assessment of a second-order probability distribution might not be easy in all situations. In the application of this model to Ellsberg's two-urn problem that was sketched above, it was assumed for convenience that the subject thought that only the two extreme cases were possible: either all red balls or all black balls in the second urn. But if the idea of a second-order probability distribution is taken seriously, then any distribution on the unit interval ought to be considered as a candidate for the second-order distribution. Perhaps a parametric family of such distributions should be considered—a beta distribution might seem like a natural choice here. But now we are building a rather elaborate model for a rather simple decision problem, which potentially requires a lot of parameter elicitation, and it is not entirely clear how such a process could be made operational, because the subjective judgment of a parameter of a second-order distribution does not directly translate into a bet on observable states of the world. However, the same type of smooth-ambiguity model can be constructed in another way that does not have so many unobservable working parts.

Suppose that the decision maker's state space can be partitioned in two or more different ways, with each partition representing a logically independent source of uncertainty. (Logical independence does not necessarily mean probabilistic independence—it just means that all combinations are possible.) In the example of Ellsberg's 2-urn problem, there is a natural partitioning of the four states into $\{\text{Red 1, Black 1}\} \times \{\text{Red 2, Black 2}\}$. Now consider the following alternative way to view the subject's preferences. Suppose that she reasons by a symmetry argument that red and black are equally likely to be drawn from urn 2, just as they are equally likely to be drawn from urn 1. Nevertheless, she feels more uncomfortable about betting on the second urn because she knows less about it. These preferences can be modeled along the following lines. Let the four states be indexed by jk , where $j=1$ [2] if the ball from urn 1 is red

[black] and $k=1$ [2] if the ball from urn 2 is red [black]. The payoffs of an act x are now indexed by jk . Suppose that the subject decomposes her beliefs as follows: first she assesses an unconditional probability distribution for the result of drawing from urn 2, and then for each possible result she assesses a conditional probability distribution for the result of drawing from urn 1. Let q denote the unconditional probability distribution for urn 2 and let p_k denote the conditional distribution for urn 1 given result k from urn 2. In this case, we have $q = (\frac{1}{2}, \frac{1}{2})$ and $p_k = (\frac{1}{2}, \frac{1}{2})$, $k = 1, 2$. Now let u and v be first- and second-order utility functions, and consider the following representation of preferences over acts:

$$U(x) = \sum_{k=1}^2 q_k v \left(\sum_{j=1}^J p_{jk} u(x_{jk}) \right)$$

When the subject evaluates an act whose payoffs depend only on urn 1, i.e., a bet whose payoffs depend on j but not on k , then her vNM utility function is effectively $u(x)$ because in that case the argument of v is the same for all k . (v will merely monotonically transform the expected utility of the urn-1 bets without affecting their preference ordering.) But if the subject evaluates an act whose payoffs depend only on urn 2, i.e., a bet whose payoffs depend on k but not on j , then her vNM utility function is effectively $v(u(x))$. If v is strictly concave, then the subject will be strictly more risk averse toward bets on urn 2 than on urn 1. Structurally, this model is a special case of the second-order probability model presented above, but it has a different, and more direct interpretation, namely that the subject is simply more averse toward betting on urn 2 than on urn 1, even if she has no reason to believe that one color is more likely than the other. This way of thinking seems very natural, and it calls attention to a very restrictive feature of the SEU model, namely that *it requires the same risk attitude to be applied to bets on all events*, regardless of the nature of those events or the quality of the information that the decision maker has on which to base her probability judgments. Remember that the vNM utility function is not supposed to represent the subjective value of consequences. Rather, it is an index of choice under risk whose role is only to model risk attitudes. What's wrong with different risk attitudes for different classes of risks?

This method of explanation of the Ellsberg and Allais paradoxes is fundamentally different from Schmeidler's and Tversky-Kahneman's since it does not depend on the existence of kinks in the indifference curves nor does it really require the assumption of state-independent utility for money nor any references to objective lotteries over prizes. (Uncertainty aversion in Schmeidler's model is a "first order" effect that is observed even with infinitesimal stakes, while in the smooth-ambiguity model it is a "second order" effect that is observed only when stakes get large enough for curvature of the utility functions to come into play.) Moreover, the characterization of risk and uncertainty aversion in this framework does not depend on prior notions of risk neutral behavior, riskless states of wealth, or probabilistic sophistication. Rather, a decision maker is risk averse if her state-preferences are convex, and she is uncertainty averse if she is uniformly more risk averse toward some classes of gambles than toward others.

This approach to modeling choice under risk and uncertainty is consistent with a view that decision makers perceive the world in terms of somewhat-distinct "frames" or "mental accounts," and that different risk attitudes—or different preference models altogether—may be appropriate for different frames. One frame might be "recreational gambling", with subframes for different kinds of games (including psychology experiments). Another frame might be

“consumer spending”, which would encompass decisions about how to choose a car insurance plan or whether to spend \$5000 on a vacation trip that might or might not live up to expectations. Other frames might be used for “risks to life and health,” or “investing for retirement,” and so on. It is often pointed out that any amount of risk aversion toward low-stakes gambles or minor casualty losses ought to make someone pathologically risk averse with respect to investments of retirement funds, yet it doesn’t, and this is not usually considered to be irrational. Rather, it suggests that *boundedly* rational decision makers find it convenient to decompose their global optimization problems into local optimization problems in which they can apply social norms and rules-of-thumb and in which an array of reasonable choices has been efficiently pre-determined from the preferences of other similar individuals in competitive markets.

Conclusions

What can we conclude from these excursions into non-expected utility theory? First of all, we can conclude that there are many ways for a decision maker to behave “rationally” without obeying the independence axiom and without having well-defined subjective probabilities. As long as her preferences under risk and uncertainty do not violate monotonicity or stochastic dominance (which they do not in the models discussed above), she does not expose herself to exploitation—at least in one-shot decisions. Many standard economic models (e.g., models of risk averse behavior and investment in securities markets) are still valid under these weaker assumptions concerning the behavior of individuals. Notice that all of the above models allow—if not require—the individual to behave like an expected utility maximizer in some choice situations, namely situations that involve only small changes in payoffs or probabilities or situations that involve only comparisons among comonotonic or sign-comonotonic acts. Thus, even if we reject the idea that the subjective expected utility is valid as a global model of rational behavior, it may still be convenient to use it as a local approximation in many situations. Among the models that relax the independence axiom while hanging on to the assumption of complete preferences, my own preference is for the general state-preference model since it does not require a unique separation of beliefs from values, does not require the decision maker to be straddling a kink in an indifference curve, and does not get into the quagmire of “consequences”. (Of course, it would be more realistic to start by relaxing the completeness assumption...)

Supplementary readings (available in back issues of *Econometrica* etc.):

If you have never read Tversky and Kahneman's original (1979) paper on prospect theory, you probably should probably do so at some point. The actual model presented in the original paper is rather crude—it refers to only 2-outcome gambles. Later work by the authors (and others) extended it to more general settings. Their 1992 paper in *Journal of Risk and Uncertainty* provides a reformulation of prospect theory in terms of cumulative probabilities, as discussed in the notes above, and.

The papers by Machina, Yaari, and Schmeidler are all milestones in the development of non-expected utility theory. Some of the technical details are messy, but the introductory sections of Machina's paper and the discussion sections of Yaari's and Schmeidler's papers give a nice perspective on the aims of theory and on the similarities and differences among different modeling approaches. Note that Machina's and Yaari's models are technically models of decision making under *risk* (i.e., objectively known probabilities) in situations involving monetary (or otherwise real-valued) payoffs. Schmeidler's model is more general insofar as it refers to more general Anscombe-Aumann acts (i.e., horse lotteries) and a subjective concept of uncertainty. Also, Schmeidler's paper focuses on the Ellsberg paradox rather than the Allais paradox.

Machina's 1989 paper on "dynamic consistency" probes an issue that, at one time, appeared to be the Achilles' heel of non-expected-utility theory: it seemingly leads to money pumps or Dutch books in decision problems involving the sequential resolution of uncertainty. Machina emphasizes that an essential feature of non-expected-utility theory is that *preferences are not separable across mutually exclusive events* as they are in standard expected utility theory. Hence it is not possible to use familiar "consequentialist" reasoning and backward-induction arguments to solve decision tree problems. Machina suggests that, for this reason, non-expected-utility maximizers will refuse to play by the rules that others might try to use to turn them into money pumps. The "mom" example is priceless. By the way, you may notice that Machina makes frequent comparisons between non-expected-utility theory and consumer theory. Indeed, his emphasis on indifference curves in triangle diagrams resembles the indifference-curve analysis used by the marginalist and ordinalist economists as well as later state-preference theorists. As such, Machina's work springs from an older tradition in economic theory, which does not give a central role to uncertainty, rather than from the von Neumann-Morgenstern framework. The idea of nonseparable preferences is quite natural in the context of consumer theory, as we have seen.