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

Criticism of subjective expected utility theory

Primary readings:

1a. Letters between Robert Aumann and Leonard Savage, 1971 (reprinted in *Essays on Economic Decisions under Uncertainty* by Jacques Drèze, 1987).

1b. "Savage Revisited" by Glenn Shafer (originally published in *Statistical Science* 1986, reprinted in *Decision Making: Prescriptive, Normative, and Descriptive Interaction*, edited by David Bell, Howard Raiffa, and Amos Tverksy, 1988)

1c. Self-test for violations of the independence axiom by Peter Wakker (from his paper "Justifying Bayesianism by dynamic decision principles")

Supplementary readings:

2a. "Risk, ambiguity, and the Savage axioms" by Daniel Ellsberg (*Quarterly Journal of Economics*, 1961; reprinted in *Decision, Probability, and Utility*, edited by Peter Gardenfors and Nils-Eric Sahlin—and many other places)

2b. "Difficulties in the theory of personal probability" by Leonard Savage and "Slightly more realistic personal probability" by Ian Hacking (panel discussion, *Philosophy of Science*, 1967)

2c. "State-dependent utility" by Mark Schervish, Teddy Seidenfeld, and Jay Kadane (*Journal of the American Statistical Association 1990*)

2d. "Separating probability elicitation from utilities" by Jay Kadane and Bob Winkler (*Journal of the American Statistical Association 1988*)

2e. "Temporal Coherence" by Michael Goldstein (in *Bayesian Statistics 2*, edited by Jose Bernardo, Morris DeGroot, Dennis Lindley, and Adrian Smith, 1985)

Guide to the readings:

Okay, now we take off the gloves for the first time (not for the last). Over the past few weeks, we have studied the development and application of the theories of subjective probability and expected utility. These theories of rational beliefs and rational preferences were synthesized by Savage into the theory of *subjective expected utility* which provides the foundation for Bayesian statistical inference, decision analysis, the economics of uncertainty, and game theory under conditions of incomplete information. Savage's theory has been presented as a *normative* theory of behavior under uncertainty to which ideally-rational decision makers ought to conform, and which social theorists can use as a modeling tool to explore the implications of the hypothesis of rationality in settings involving uncertainty. We have also looked at a number of models of non-expected utility that extend the EU and SEU models. This week's readings delve more deeply into some very serious problems in the theory--problems that that run deeper than the inability of mere mortals to behave with the perfect consistency that the theory demands. The problems concern the very nature of consequences, acts, preferences, and beliefs that are the primitives in the theory. The primary readings are largely non-mathematical and are eloquently written, so I will not go into much detail trying to summarize their main arguments—they will reward a close reading.

1a. Letters between Robert Aumann and Leonard Savage, 1971 (reprinted in *Essays on Economic Decisions under Uncertainty* by Jacques Drèze, 1987).

> This wonderful exchange of letters calls attention to a profound and under-appreciated problem in the theory of subjective expected utility: the difficulty of separating probabilities from utilities. (The same problem is explored in much more depth in the following paper by Shafer.) Aumann was well aware of the problem from his own work on axiomatizing SEU theory: he and Anscombe had developed a simpler set of axioms in 1963, in which the objects of choice are "horse lotteries" that map states of nature to objective lotteries over consequences. The difficulty arises in situations where it is impossible to define "consequences" in such a way that they can be expected to have utilities that are state-independent. Aumann presents several poignant examples of such situations. Judge for yourself whether Savage's response is persuasive! These letters were written in January 1971; tragically, Savage died in November of the same year at the age of 54. (By the way, Drèze is another of the major figures in the development of the theory of markets under uncertainty, and he has also done pioneering work on the subject of state-dependent utility. This book of essays is a nice survey of some of his contributions.)

1b. "Savage Revisited" by Glenn Shafer (originally published with discussion in *Statistical Science* 1986, reprinted without the discussion in *Decision Making: Prescriptive, Normative, and Descriptive Interaction*, edited by David Bell, Howard Raiffa, and Amos Tverksy, 1988)

This is an outstanding paper that both explains and deconstructs the Savage axioms. Shafer very carefully discusses all the assumptions in the theory and walks through a detailed analysis of the major axioms. He presents a thorough discussion of the often-criticized independence axiom (P2), citing the famous paradoxes of Allais and Ellberg, but also points out that there are more important problems with the axioms that separate probability from utility (P3 and P4). Shafer refers to P3 as the assumption that "value can be purged of belief" and to P4 as the assumption that "belief can be discovered from preference." They are the same assumptions that lie at the heart of the problems pointed out in Aumann's letter. Shafer questions whether it is possible *even in principle* to shrink a "grand world" decision problem into a "small world" in which these axioms would be true. The underlying problem is one of defining "states of the person" (consequences) that are sufficiently fine-grained and yet somehow logically distinct from "states of nature."

Shafer also takes aim at the assumption that *preferences* are a psychologically primitive relation. He argues that preferences are deliberately *constructed* as the need arises—by starting from probability and value judgements, not the other way around—and that a decision maker often will have better things to do with her time than construct preferences for a large number of purely imaginary alternatives. (The idea that preferences are perhaps the wrong primitive on which to base rational choice theory is another major theme to which we will return later in the course. By the way, Shafer is perhaps best known for his work on the theory of *belief functions*, a non-Bayesian approach to modeling uncertainty that was popular for some time in the artificial intelligence community. He has also written extensively on the constructive nature of beliefs and preferences.)

The accompanying discussion by Lindley, Dawid, Dawes, Fishburn, and Pratt is also very much worth reading. Pratt's comments are particularly caustic: he characterizes Shafer's paper as "a regressive exercise in overliteral misreading and straw battle." Draw your own conclusions!

Note: if you have trouble figuring out Shafer's small-world analysis of Savage's omelette problem at the end of the paper, see the very helpful decision-tree solution of the same problem in Phil Dawid's comment. Dawid shows the decision tree for the grand-world version of the problem. The key thing to note here is that in the corresponding small-world version of the tree, the expected utilities on the three main decision branches (i.e., the three concrete acts) must be 14, 20, and 13, the same as in the grand-world tree. (This is because preferences among acts in the small world are assumed to agree exactly with their grand-

world counterparts.) However, the small-world probability of a good 6th egg and the small-word utility of a five-egg omelet (at nodes 8, 9, and 10) and the small-world utility of a six-egg omelet (at nodes 5 and 7) remain to be determined. By solving three equations in these three unknowns, using the known expected utilities on the decision branches, you obtain p(good) = 7/13, u(5-egg omelet) = 13, and u(6-egg omelet) = 26. The key thing, as Dawid points out, is that a 5-egg omelette is a final consequence in the small world, so it must have the same utility everywhere it appears, whereas in the grand world a 5-egg omelette has different *expected* utilities depending on whether the 6th egg was good or rotten. The probability of "good" in the small world must be distorted to correct for this difference.

1c. Self-test for violations of the independence axiom by Peter Wakker (from his paper "Justifying Bayesianism by dynamic decision principles"**)**

Do *you* violate the independence axiom. If so, where and why? This exercise will help you find out. Wakker has drawn different decision tree diagrams (i.e., "extensive-form" representations) of the Allais problem, with payoffs scaled up to adjust for inflation. Each tree represents a "protocol" in which decisions and events happen in a particular sequence, and neighboring trees differ by a single transformation operation which may or may not make a difference to you. Under subjective expected utility theory, all the trees are essentially equivalent, so if you are an SEU-maximizer you should not reverse your preferred decision (up or down) as you move from (a) to (e). Under various alternative theories they need *not* be equivalent. Wakker's diagram shows which authors have relaxed which equivalence assumptions. (Note that a single "normal form" decision problem—i.e., a tabular problem as given in my notes above—may have different extensive form representations. One of the convenient implications of the independence axiom is that it implies an equivalence between normal and extensive forms, so the different extensive forms should have the same solution. When independence is weakened, the different extensive forms may have different solutions.)

Supplementary readings: The following papers provide additional technical analysis of the issues raised in the other papers.

2a. "Risk, ambiguity, and the Savage axioms" by Daniel Ellsberg (*Quarterly Journal of Economics*, 1961)

This classic paper calls attention to another apparent problem with the subjective theory of probability: behavior that suggests aversion to "ambiguity" in the probabilities attached to events. If you always prefer to bet on events whose probabilities are objectively known rather than similar events whose probabilities are only subjectively determined, "you are now in trouble with the Savage axioms." (Savage's axioms imply that once you have determined your subjective probabilities, you should use them to make choices exactly as if they were objective probabilities.) Ellsberg conducted his original study under "absolutely nonexperimental conditions" using eminent decision theorists (Marschak, Samuelson, Raiffa, Debreu, et al.) as subjects, although his results have been widely replicated, and "Ellsberg urn" problems continue to be a source of inspiration for theories of non-expected utility, which we will discuss in the next class.

2b. "Difficulties in the theory of personal probability" by Leonard Savage and "Slightly more realistic personal probability" by Ian Hacking

These two papers were part of a panel discussion published in *Philosophy of Science* in 1967. Savage's paper reflects his "mature views" as mentioned by Fishburn in his comment on Shafer's paper. Hacking's paper very importantly challenges another assumption that is central to information economics, namely the assumption that Bayes' theorem is a model of *learning over time*. (Goldstein's paper in the supplementary readings also addresses this issue.) Recall that Bayes' theorem states that $P(H|E) \propto P(H)P(E|H)$. Now, formally, this is just a relation that logically must be satisfied by certain conditional and unconditional probabilities. But in applications, where H denotes a hypothesis and E denotes the outcome of a statistical experiment, it is common to refer to refer to $p(H)$ as your belief in the hypothesis "prior" to performing the experiment and $p(H|E)$ as your belief in hypothesis "posterior" to performing the experiment, where "posterior" is often casually interpreted to mean *posterior in time*. Under this interpretation of P(H|E), models of "Bayesian learning" are often constructed in which agents update their beliefs over time by repeated, mechanical applications of Bayes' theorem to data that they observe. In fact, there is nothing in Savage's or de Finetti's axioms that justifies this interpretation, as this paper points out, and "Bayesian learning" is an oxymoron.

"Well known properties of [Bayes'] theorem lead us to a model of learning from experience.... The idea of the model of learning is that P(H|E) represents one's personal probability after one learns E.

But formally the conditional probability represents no such thing. If, as in all of Savage's work, conditional probability is a defined notion, then P(H|E) stands merely for the quotient of two probabilities. It in no way represents what I have learned after I take E as a new datum point....

A man knowing E would be incoherent if the [betting] rates offered on H unconditionally differed from his rates on H conditional on E. But no incoherence obtains when we shift from the point before E is known to the point after it is known.

The man's P(H|E) before learning E differs from his P(H|E) after learning E. Why not, he says: the change represents how I have learned from E!"

2c. "State-dependent utility" by Mark Schervish, Teddy Seidenfeld, and Jay Kadane (*Journal of the American Statistical Association 1990*)

This paper discusses the problem of the state-dependence of utilities and its implications for the assessment of subjective probabilities from observations of preference. It presents both Savage's and Anscombe-Aumann's axioms for subjective expected utility, and shows that they both require the assumption of a "constant" of utility across states in order to separate probability from utility. The paper also gives a very nice example of how the same problem arises in de Finetti's method of probability elicitation: suppose that you assess someone's probabilities for states of nature by observing the prices a person is willing to pay for lottery tickets on events, where the payoffs of the bets are expressed in dollars. Now suppose you repeat the same exercise with the payoffs expressed in another currency, say, Japanese yen. Should you get the same results? Maybe or maybe not! The uncertain state might be tomorrow's exchange rate between the dollar and the yen. If this is the case, it is reasonable that the relative prices should be different. But which reveals the person's "true" probability? That depends on what counts as a "constant" of utility! If the person buys consumption goods in dollars, then perhaps dollars yield constant utility; if she consumes in yen, then perhaps yen yield constant utility. In general, how do we know what is the right constant? Schervish, Seidenfeld, and Kadane suggest that this problem is potentially a "catastrophe" for the theory of Bayesian statistical inference, which assumes that it is meaningful to speak of probabilities separately from utilities, and they point to several approaches that have been suggested for solving this problem. (As for myself, I think the problem of operationally measuring "true" probability is fundamentally insoluble—but not to worry, it is not catastrophic either. It just requires a decision theory that does not depend heavily on the separation of belief from value.)

2d. "Separating probability elicitation from utilities" by Jay Kadane and Bob Winkler (*Journal of the American Statistical Association 1988*)

This paper explores the same problem in the context of de Finetti's two money-based methods of probability elicitation: betting rates and scoring rules. Even if an individual's "true" utility function is state-independent, her marginal utility for any currency can be state-dependent if she has a nonlinear utility function and significant *prior stakes* in the outcomes of events. In this case, the probabilities revealed by money-based elicitation schemes will be distorted by the marginal utilities. As we discussed last week, the individual's betting rates will reveal her "risk neutral" probabilities rather than her "true" probabilities. But is this a problem for Bayesian theory…?

2e. "Temporal Coherence" by Michael Goldstein (in *Bayesian Statistics 2*, edited by Jose Bernardo, Morris DeGroot, Dennis Lindley, and Adrian Smith, 1985)

This paper shows that the "dynamic assumption" criticized by Hacking actually can be modeled within a de Finetti-type framework, but it requires a formal definition of *coherence over time*. To construct this, you first need to treat your own future probabilities as uncertain quantities about which you can form beliefs--and on which you can bet--today. Temporal coherence (i.e., the avoidance of Dutch books under these conditions) requires that the prior probabilities that you hold today must equal your expected value of the posterior probabilities you will hold tomorrow--but tomorrow's probabilities may be random to some extent. Thus, at any point in time, you must expect your beliefs to unfold over time according to a martingale process. This is essentially the same martingale property that characterizes arbitrage-free stock and option prices in a securities market. So, whereas the concept of "Bayesian learning" is often used as a microfoundation for modeling the evolution of prices in stock markets, it is perhaps more appropriate to take the stock market as a model of the intrinsic volatility of the beliefs that exist in the mind of an individual.

> *"The question as to what information is conveyed by the expression P(H|E) is central to the Bayesian argument. Most developments of Bayesian theory proceed by defining conditional probabilities and deriving their properties in terms of called off penalties or bets, determined before the actual conditioning event is revealed. This definition is then taken to represent your actual beliefs about H having seen E. This transition is not justified and is rarely even pointed out. However, it is clearly a false transition on both theoretical and practical grounds.*
>
> *As no coherence principles are used to justify the equivalence of conditional and a posteriori probabilities, this assumption is an arbitrary imposition on the subjective theory. As Bayesians rarely make a simple*

updating of actual prior probabilities to the corresponding conditional probabilities, this assumption misrepresents Bayesian practice."