

Chapter 2 of Donald Meyer (ed.), *The Economics of Risk*.
W.E. Upjohn Institute for Employment Research, 2003

States of the World and the State of Decision Theory

Mark J. Machina

Department of Economics
University of California, San Diego
La Jolla, California 92093-0508

Werner Sichel Public Lecture on the Economics of Risk, presented at Western Michigan University on November 7, 2001 and co-sponsored by the W. E. Upjohn Institute for Employment Research. I would like to thank Ted Groves, Donald Meyer and Joel Sobel for helpful comments. This material is based upon work supported by the National Science Foundation under Grant No. 9870894. All opinions and errors are my own.

copyright: W.E. Upjohn Institute for Employment Research

STATES OF THE WORLD AND THE STATE OF DECISION THEORY

TWO EXTRAORDINARY NON-SCIENTISTS

SCIENTIFIC CONSIDERATIONS IN THE THEORY OF CHOICE

Scientific Modelling “From the Outside In”

Issues of Observability

Issues of Classification

Issues of Measurability

CHOICE UNDER OBJECTIVE UNCERTAINTY

Outcomes, Probabilities and Objective Lotteries

The Expected Utility Hypothesis

Violations of the Expected Utility Hypothesis

Non-Expected Utility Models of Risk Preferences

CHOICE UNDER SUBJECTIVE UNCERTAINTY

States, Events, Outcomes and Acts

The Hypothesis of Probabilistic Sophistication

Violations of the Hypothesis of Probabilistic Sophistication

Non-Probabilistically Sophisticated Models of Risk Preferences and Beliefs

ARE STATES OF NATURE PRESCRIPTIVE OR DESCRIPTIVE?

Criteria for Normative Applications

Do Decision-Makers See the State Space or Do They Construct the State Space?

ISSUES OF OBSERVABILITY, CLASSIFICATION AND MEASUREMENT

Independent Observability and the Exogeneity of States

Ex Ante Observability versus Ex Post Observability of States

Issues of Classification and Measurement

ISSUES OF STRUCTURE: ALMOST-OBJECTIVE UNCERTAINTY

Properties of Purely Objective Events

Almost-Equally-Likely Events and Almost-Fair Bets

Almost-Objective Events, Acts and Mixtures

Beliefs and Betting Preferences over Almost-Objective Events

APPENDIX 2A: PROPERTIES OF THE SMOOTH FUNCTION APPROACH TO NON-PROBABILISTICALLY SOPHISTICATED MODELS

APPENDIX 2B: ALMOST-OBJECTIVE UNCERTAINTY

REFERENCES

2

States of the World and the State of Decision Theory

Mark J. Machina
University of California, San Diego

TWO EXTRAORDINARY NONSCIENTISTS

Almost 20 years ago, I briefly knew a man by the name of Craig. Although he died about a year after I met him, I've thought about him ever since. Craig had this uncanny ability to converse with a person for a few minutes, and then announce what make and model of car they drove. Neither I, nor anyone I ever spoke to, had ever seen him get it wrong. Craig was never able to explain how he did it, and his unique ability followed him to the grave.

What Craig had perfected was an impressive skill—perhaps even an art—but it was not science. It was not science because it was not a procedure that he could verbally communicate or write down, so that other people in other places or other times could do it also. One of the defining features of scientific activity is that it generates a body of knowledge and techniques that *can* be communicated and utilized by others in this way.

I also knew a woman named Tula with an equally impressive ability. Tula was able to predict how well a person's day would go, based on the shape, size, and color of the *aura* they emitted in the morning. And in contrast to Craig, she could even explain the specifics of her method. For example, if your aura was round and blue, you would have good luck all day. But if it was square and yellow, then you'd best go back home and stay in bed. Tula had prepared a chart with the complete relationship between properties of your aura and the upcoming features of your day, so if you had a copy of the chart, you just needed a daily reading of your aura. Although Tula's success rate wasn't per-

fect (like Craig's was), it still compared favorably with standard medical, meteorological, and macroeconomic predictions, and most of her friends would stop by each morning for a quick reading of their aura, and then go away to consult their chart.

By constructing and distributing her chart, Tula had codified and communicated features of her technique in a way that Craig never could. But since Tula was the only one who could *see* these auras, what she was doing still was not science. An activity is not science unless it involves techniques that others can also apply *as well as* variables that others can observe.

The purpose of this chapter is to examine one of the most important theoretical constructs of modern decision theory—namely, the concept of *states of the world* or *states of nature*—from the point of view of these and similar scientific considerations. Are states of nature inherently descriptive or prescriptive objects? Do individuals making choices under uncertainty *face* these states of nature, or do they *create* them? Are states external and independently observable, like an individual's commodity demand levels, or are they internal and not directly observable, like utility or marginal utility levels? In addressing these questions, I will offer an overview of how researchers have sought to represent the concept of uncertainty, from the original formulation of probabilities and "objective uncertainty" in the seventeenth century, through Leonard Savage's twentieth century formulation of states of nature and "subjective uncertainty," to current work which seeks to eliminate—or at least redefine—the distinction between objective and subjective uncertainty. The following section presents some scientific issues common to all theories of choice, whether under certainty or uncertainty. The next two sections sketch out the current theories of choice under objective and subjective uncertainty. After that, I address the question of whether states of nature should be considered descriptive or prescriptive constructs, and then I consider scientific issues related to the observability and measurement of states of nature. The final section concludes with current work on the relationship between subjective and objective uncertainty.

SCIENTIFIC CONSIDERATIONS IN THE THEORY OF CHOICE

Scientific Modeling “From the Outside In”

The human decision-making process may well be one of the most complicated systematic phenomena in the universe. In terms of the point of view of the scientific observer, it is certainly unique. On the one hand, a scientist trying to model this process is like an anatomist in the days before anesthesia and vivisection—scientists can observe and to some extent even control external influences on a system, and can observe the resulting behavior of the system as a whole, but they cannot “get inside” to observe its constituent parts at work. On the other hand, every scientist is a human decision maker with powers of self-consciousness and self-reflection. However, self-reflection of our decision-making processes has not produced that much more “hard science” than has, say, self-reflection of our breathing or digestive processes.

While advances in neuroscience may ultimately do for decision theory what vivisection did for anatomy, decision theory currently remains very much a “black box” science. Although decision theorists can (and do) use introspection to suggest theories and hypotheses, the rigorous science consists of specifying mutually observable independent variables (in particular, the objects of choice available for selection), mutually observable dependent variables (the selected alternative), and refutable hypotheses linking the two. In other words, choice theory attempts to explain why particular alternatives are selected from a set of available choices.

Issues of Observability

Because decision scientists cannot perform dissection, they are subject to a greater scientific discipline than that required of anatomists. If a decision scientist tried to account for an individual’s purchases of bananas as the direct result of something like an “appetite for fruit,” we would not know how to test this hypothesis—that is, we would not know how to independently “look for” such an appetite, even if we had a scalpel and an open, anesthetized brain. Such unob-

servable constructs like appetites, utility, and preferences can—and do—play a role in scientific decision theory, but only as inside links in a causal chain that ultimately starts with fully observable independent variables and ultimately ends with fully observable dependent variables. For example, given the joint hypothesis that well-defined commodity preferences exist and are also stable from day to day, standard consumer theory allows us to infer enough information about these preferences from an individual's past demand behavior to be able to make refutable predictions about their future demand behavior, even for some combinations of prices and income never before observed.

In the following sections, we shall see that in passing from choice over certain commodity bundles to choice over uncertain prospects (either “objective lotteries” or “subjective acts”), hypotheses involving the unobservable constructs of commodity preferences and utility functions can be replaced by hypotheses involving the unobservable constructs of risk preferences and beliefs, which also link observable independent to observable dependent variables. Whether the notion of “states of nature” can similarly serve remains to be discussed.

Issues of Classification

In order for a variable or phenomenon to satisfy the criterion of “scientific observability,” it is not enough that more than one scientist be able to see it—it is not even enough that a camera be able to record it. Rather, a variable is only scientifically observable if independent observers can agree on their description of what they have just observed. Thus, while a scientist can photograph facial expressions, they cannot be said to have photographed expressions of emotion unless there is a well-defined specification of which expressions correspond to each emotion, and independent observers predominantly agree in their assignment of emotions to each photograph. In other words, scientific observability requires well-defined and commonly accepted classification schemes for the observations, sufficient for grouping and comparing such observations, and relating them to general hypotheses and theories.

Just as different types of variables can have different degrees of observability, different classification schemes will have different degrees of common agreement. Thus, in regular consumer theory, we

are much more prone to classify commodities and define preferences in terms of category schemes like {"fruits," "vegetables," "grains"} compared to schemes like {"delicious foods," "filling foods," "unpleasant foods"}. Although the latter scheme is in some sense much more directly connected to any given individual's preferences than the former scheme, the latter scheme cannot be defined independently of the particular consumer being studied. Since foods cannot be classified according to this latter scheme prior to observation of the consumer's (verbal or choice) behavior, it cannot be used as a classification scheme for independent variables. Categories like "delicious foods," "unpleasant foods," etc. can be defined for *dependent variables*, however, either on the basis of the consumer's verbal expressions, or on the basis of their past purchases or consumption behavior. Thus, whether a given classification scheme does or does not satisfy the criterion of scientific observability may well depend upon whether the scheme is intended to be applied to the independent variables or to the dependent variables of a theory.

Issues of Measurability

The above example of classifying facial photographs into different categories of emotions is an example of a *qualitative* classification of the basic observations. Although qualitative categories and qualitative variables are perfectly valid in the physical, biological, and social sciences, theories and hypotheses are most powerful when they involve *quantitative* independent and dependent variables. Many economists are of the opinion that economics has a more impressive scientific track record than anthropology because economists work with numerical variables such as prices, quantities, and income, rather than with qualitative variables like trust, group identification, or loyalty. Most theories and hypotheses involving quantitative independent and dependent variables are easier to test, to fine tune, and if necessary, to revise, than most theories and hypotheses involving qualitative variables.

Is uncertainty an inherently qualitative or quantitative construct? In the following sections we shall see that one of the two primary methods of representing uncertainty—the so-called "objective approach"—represents uncertainty quantitatively, via numerical probabilities. On the other hand, the other primary method—the so-called "subjective

approach”—has traditionally represented uncertainty in a qualitative manner, via an unstructured set of states of nature. However, in the final section of this paper, we see that taking a measurable, quantitative approach to subjective uncertainty can enhance its power, and in many senses can serve as an almost complete substitute for what may be considered the more ad hoc assumptions made about the world in the objective approach.

CHOICE UNDER OBJECTIVE UNCERTAINTY

Outcomes, Probabilities, and Objective Lotteries

The earliest formal representation of uncertainty came from founders of modern probability theory such as Pascal and Fermat. In this approach, the uncertainty attached to any event is represented by a numerical probability p between 0 and 1. Because probability theory derived from the study of games of chance that involved virtually identical repeated events, such probabilities were held to be intrinsic properties of the events in the sense that an object's mass is an intrinsic property of the object. These probabilities could either be calculated from the principles of combinatorics, for an event such as being dealt a royal flush, or measured by repeated observation, for an event like a bent coin landing heads up.

For an individual making a decision under objective uncertainty, the objects of choice are *objective lotteries* of the form $\mathbf{P} = (x_1, p_1; \dots; x_m, p_m)$, which yield outcome x_i with objective probability p_i , where $p_1 + \dots + p_m = 1$. The theory of choice under uncertainty treats lotteries in a manner almost identical to the way it treats commodity bundles under certainty. That is, each individual's preferences over such lotteries can be represented by a real-valued *preference function* $V(\cdot)$, in the sense that for any pair of lotteries $\mathbf{P}^* = (x_1^*, p_1^*; \dots; x_{m^*}^*, p_{m^*}^*)$ and $\mathbf{P} = (x_1, p_1; \dots; x_m, p_m)$, the individual prefers \mathbf{P}^* over \mathbf{P} if and only if $V(\mathbf{P}^*) = V(x_1^*, p_1^*; \dots; x_{m^*}^*, p_{m^*}^*)$ exceeds $V(\mathbf{P}) = V(x_1, p_1; \dots; x_m, p_m)$, and is indifferent between the two lotteries if and only if $V(\mathbf{P}^*) = (x_1^*, p_1^*; \dots; x_{m^*}^*, p_{m^*}^*)$ exactly equals $V(\mathbf{P}) = V(x_1, p_1; \dots; x_m, p_m)$.¹

The Expected Utility Hypothesis

In standard consumer theory, the preference function over commodity bundles is typically assumed to have certain mathematical properties but is typically *not* hypothesized to take any specific functional form, such as the Cobb-Douglas or Constant Elasticity of Substitution form. Specific functional forms are typically only used when absolutely necessary, such as in empirical estimation, calibration, or testing.

In contrast, the standard theory of choice under objective uncertainty typically *does* assume (or does assume axioms sufficient to imply) a specific functional form for the individual's preference function over lotteries, namely the objective expected utility form $V_{EU}(x_1, p_1; \dots; x_m, p_m) = U(x_1) \cdot p_1 + \dots + U(x_m) \cdot p_m$ for some von Neumann-Morgenstern utility function $U(\cdot)$. Mathematically, the characteristic features of this functional form are that it is additively separable in the distinct (x_i, p_i) pairs, and also that it is linear in the probabilities. The term "expected utility" arises since it can be thought of as the mathematical expectation of the variable $U(x)$ (the individual's "utility of wealth") if wealth x has distribution $\mathbf{P} = (x_1, p_1; \dots; x_m, p_m)$. The literature on choice under uncertainty has generated a number of theoretical results linking the shape of the utility function to aspects of the individual's attitudes toward risk, such as risk aversion or comparative risk aversion for a pair of individuals. Excellent discussions of the foundations and applications of expected utility theory can be found in standard graduate level microeconomic texts such as Kreps (1990, Chapter 3), Mas-Colell, Whinston, and Green (1995, Chapter 6), and Varian (1992, Chapter 11).

Violations of the Expected Utility Hypothesis

Although the expected utility model is sometimes viewed as being quite flexible (since the von Neumann-Morgenstern utility function could have any shape), it does generate refutable predictions. Unfortunately, there is a growing body of evidence to suggest that individuals' preferences over lotteries tend to systematically violate some of these predictions. Risk preferences tend to systematically depart from the expected utility property of linearity in the probabilities. The most

notable example of this is the well-known Allais Paradox (Allais 1953), which asks individuals to rank each of the following pairs of lotteries (where \$1M denotes \$1,000,000):

$$a_1 : \left\{ \begin{array}{l} 1.00 \text{ chance of } \$1\text{M} \\ 0.00 \text{ chance of } \$0 \end{array} \right. \quad \text{versus} \quad a_2 : \left\{ \begin{array}{l} 0.10 \text{ chance of } \$5\text{M} \\ 0.89 \text{ chance of } \$1\text{M} \\ 0.01 \text{ chance of } \$0 \end{array} \right.$$

$$a_3 : \left\{ \begin{array}{l} 0.10 \text{ chance of } \$5\text{M} \\ 0.90 \text{ chance of } \$0 \end{array} \right. \quad \text{versus} \quad a_4 : \left\{ \begin{array}{l} 0.11 \text{ chance of } \$1\text{M} \\ 0.89 \text{ chance of } \$0 \end{array} \right.$$

Experiments by Allais and others have found that the modal (and in some studies, the majority) choices are for a_1 over a_2 in the first pair, and a_3 over a_4 in the second pair. However, a preference for a_1 in the first pair implies that the utility function satisfies the inequality $0.11 \cdot U(\$1\text{M}) > 0.10 \cdot U(\$5\text{M}) + 0.01 \cdot U(\$0)$, whereas a preference for a_3 in the second pair implies $0.11 \cdot U(\$1\text{M}) < 0.10 \cdot U(\$5\text{M}) + 0.01 \cdot U(\$0)$, which is a contradiction.

Although the Allais Paradox was originally dismissed as an isolated example, subsequent work by MacCrimmon and Larsson (1979), Kahneman and Tversky (1979), and others have uncovered a qualitatively similar pattern of departure from the expected utility hypothesis of linearity in the probabilities, over a large range of probability and payoff values (see Machina 1983, 1987 for surveys of this evidence).

NON-EXPECTED UTILITY MODELS OF RISK PREFERENCES

Responses to the above-mentioned violations of the expected utility hypothesis have taken two forms. One branch of the literature has proceeded by positing more general functional forms for the preference function (Edwards 1955, 1962; Kahneman and Tversky 1979; Chew 1983; Fishburn 1983; Quiggin 1982; and Yaari 1987). Such forms accommodate most of the observed departures from linearity in the probabilities, and, given the appropriate curvature assumptions, can

also exhibit standard features like risk aversion, comparative risk aversion, etc.

A second line of work in non-expected utility theory proceeds in a manner closer to that of standard consumer theory—rather than adopting some new functional form, it generalizes the expected utility property of linearity in the probabilities to its natural extension of smoothness in the probabilities (e.g., Machina 1982). That is, it treats the preference function $V(x_1, p_1; \dots; x_m, p_m)$ as a general smooth function, and studies how properties of its probability derivatives relate to attitudes toward risk. This approach finds that much of expected utility theory is analytically robust to departures from linearity in the probabilities.²

CHOICE UNDER SUBJECTIVE UNCERTAINTY

States, Events, Outcomes, and Acts

From a mathematical perspective, the representation of uncertainty by means of additive, numerical probabilities allows us to apply the tremendous body of analytical results of modern probability theory (e.g., Feller 1968, 1971; Billingsley 1986). But from a modeling perspective, the assumption that uncertainty comes prepackaged with well-defined, measurable “objective” probabilities is unrealistic. Outside of the gambling hall, most economic decisions and transactions involving uncertainty—investment decisions, search decisions, insurance contracts, financial instruments—are defined in terms of uncertain events rather than numerical probabilities.

This approach to representing uncertainty and uncertain prospects—formalized by Savage (1954) and now known as the subjective approach—involves the following basic constructs:

$\mathcal{X} = \{ \dots, x, \dots \}$ an arbitrary space of outcomes or consequences.

$S = \{ \dots, s, \dots \}$	a space of mutually exclusive and exhaustive states of nature, representing all possible alternative unfoldings of the world.
$E = \{ \dots, E, \dots \}$	an algebra of events, each a subset of S .
$f(\cdot) = [x_1 \text{ on } E_1; \dots; x_m \text{ on } E_m]$	subjective act yielding outcome x_i in event E_i , for some partition $\{E_1, \dots, E_m\}$ of S (or equivalently, yielding outcome $f(s)$ in state s).
$A = \{ \dots, f(\cdot), \dots \}$	the set of all such subjective acts.
$W(\cdot)$ and \succ	the individual's preference function and corresponding preference relation over A .

A wonderful example of the use of this framework to represent an uncertain decision was provided by Savage (1954, pp. 13–15): Say you are making omelets and have already broken five of your six eggs into a mixing bowl. The decision you must make is: Do you break the sixth egg? The uncertainty arises from the fact that this sixth egg has been around for some time and might be rotten. You can either break this egg into the bowl with the other eggs, break it into a separate saucer to inspect it, or throw it away unbroken. Savage represents this problem in terms of states, acts, and outcomes by means of the following table:

Act	State	
	Egg is good	Egg is rotten
Break into bowl	Six-egg omelet	No omelet, and five good eggs destroyed
Break into saucer	Six-egg omelet, and a saucer to wash	Five-egg omelet, and a saucer to wash
Throw away	Five-egg omelet, and one good egg destroyed	Five-egg omelet

The Hypothesis of Probabilistic Sophistication

Although the subjective approach drops the assumption that uncertainty is defined in terms of numerical probabilities, it still allows for

individuals to possess probabilistic beliefs, with the feature that such beliefs may now differ across individuals. Formally, an individual is said to be probabilistically sophisticated, with a subjective (or personal) probability measure $\mu(\cdot)$ over the events E , if their preference function $W(\cdot)$ over subjective acts takes the form

$$W_{PS}(f(\cdot)) = W_{PS}(x_1 \text{ on } E_1; \dots; x_m \text{ on } E_m) = V(x_1, \mu(E_1); \dots; x_m, \mu(E_m))$$

for some (not necessarily expected utility) preference function $V(\mathbf{P}) = V(x_1, p_1; \dots; x_m, p_m)$ over lotteries. That is to say, an individual is *probabilistically sophisticated* if their uncertain beliefs can be completely summarized by a subjective probability $\mu(E)$ attached to each event E , and the individual evaluates each subjective act $f(\cdot) = [x_1 \text{ on } E_1; \dots; x_m \text{ on } E_m]$ solely on the basis of its implied probability distribution $(x_1, \mu(E_1); \dots; x_m, \mu(E_m))$ over outcomes. This representation of $W_{PS}(\cdot)$ as the composition of a preference function $V(\cdot)$ over lotteries and a subjective probability measure $\mu(\cdot)$ over events is now referred to as the classical separation of risk preferences from beliefs.

Violations of the Hypothesis of Probabilistic Sophistication

Savage's (1954) joint axiomatization of expected utility risk preferences and probabilistic beliefs, employing an expected utility function for the risk preference function, has been justly termed "the crowning glory of choice theory" (Kreps 1988, p.120). However, the violations of expected utility first observed by Allais were soon matched by violations of probabilistic sophistication, even in situations involving the simplest forms of subjective uncertainty. The most famous of these examples, known as the Ellsberg Paradox (Ellsberg 1961, 2001), involves drawing a ball from an urn containing 30 red balls and 60 black or yellow balls in an unknown proportion. The following table illustrates four subjective acts defined over the color of the drawn ball, when the entries in the table are payoffs or outcomes:

	30 balls		60 balls	
	red	black	black	yellow
$f_1(\cdot)$	\$100	\$0	\$0	\$0
$f_2(\cdot)$	\$0	\$100	\$100	\$0
$f_3(\cdot)$	\$100	\$0	\$0	\$100
$f_4(\cdot)$	\$0	\$100	\$100	\$100

When faced with these choices, most subjects prefer act $f_1(\cdot)$ over $f_2(\cdot)$, on the grounds that the probability of winning \$100 in $f_1(\cdot)$ is guaranteed to be 1/3, whereas in $f_2(\cdot)$ it could range anywhere from 0 to 2/3. Similarly, most subjects prefer $f_4(\cdot)$ over $f_3(\cdot)$, on the grounds that the probability of winning \$100 in $f_4(\cdot)$ is guaranteed to be 2/3, whereas in $f_3(\cdot)$ it could range anywhere from 1/3 to 1. Although this reasoning may well be sound, it is inconsistent with the hypothesis of probabilistic beliefs. That is, there is *no* triple of subjective probabilities $\{\mu(\text{red}), \mu(\text{black}), \mu(\text{yellow})\}$ that can simultaneously generate a preference for $f_1(\cdot)$ over $f_2(\cdot)$ and for $f_4(\cdot)$ over $f_3(\cdot)$, since a probabilistically sophisticated individual would only exhibit the former ranking when $\mu(\text{red}) > \mu(\text{black})$, and only exhibit the latter ranking when $\mu(\text{red}) < \mu(\text{black})$.

Ellsberg also presented what many feel to be an even more fatal example, involving two urns:

	50 balls		50 balls		100 balls	
	red	black	black	red	red	black
$g_1(\cdot)$	\$100	\$0	\$0	\$100	\$100	\$0
$g_2(\cdot)$	\$0	\$100	\$100	\$0	\$0	\$100
$g_3(\cdot)$					\$100	\$0
$g_4(\cdot)$					\$0	\$100

In this example, most subjects are indifferent between $g_1(\cdot)$ and $g_2(\cdot)$, are indifferent between $g_3(\cdot)$ and $g_4(\cdot)$, but strictly prefer either of $g_1(\cdot)$ or $g_2(\cdot)$ to either of $g_3(\cdot)$ or $g_4(\cdot)$. It is straightforward to verify that there exist no pair of subjective probabilities $\{\mu(\text{red}), \mu(\text{black})\}$ for the right-hand urn—50:50 or otherwise—that can generate this set of preference rankings. Such examples illustrate the fact that in situations (even simple situations) where some events come with probabilistic information

and some events (termed *ambiguous events*) do not, subjective probabilities do not always suffice to fully encode all aspects of an individual's uncertain beliefs. Since most real-world events do not come with such probabilistic information, Ellsberg's Paradoxes and related phenomenon deal a serious blow to the hypothesis of probabilistic sophistication.

Non-probabilistically Sophisticated Models of Risk Preferences and Beliefs

Just as the Allais Paradox and similar evidence led to the development of non-expected utility models of risk preferences, Ellsberg's Paradoxes and similar phenomena have inspired the development of non-probabilistic models of preferences over subjectively uncertain acts. Such work has also progressed along two lines. One line replaces the subjective expected utility function with more general functional forms.³

The second line of research on non-probabilistic models treats $W(x_1 \text{ on } E_1; \dots; x_m \text{ on } E_m)$ as a general smooth function of the events E_1, \dots, E_m , and show how properties of $W(\cdot)$'s *event-derivatives* relate to features of both beliefs and attitudes toward risk, again taking expected utility as its base case. Appendix 2A presents mathematical features of this line of research.

ARE STATES OF NATURE PRESCRIPTIVE OR DESCRIPTIVE?

The second section in this chapter argued that the scientific suitability of a particular theoretical construct—in that case it was a particular classification scheme for food—could depend on whether the construct was meant to be applied to the independent variables of a theory or its dependent variables. This section addresses a similar issue, namely that certain criteria for suitable specification of the states of nature can depend upon whether the states are to be used for positive (that is, descriptive) versus normative (that is, prescriptive) purposes.

Since its inception, expected utility theory has always straddled the boundary between being a descriptive and a prescriptive model of decision making under uncertainty. Even its original presentation by Bernoulli (1738) as a “solution” to the St. Petersburg Paradox can be alternatively interpreted as either a description of why people *don't* assign an infinite certainty equivalent to the Petersburg Game, or a prescription for why an individual *shouldn't* assign an infinite certainty equivalent to the game. Two centuries later, proponents of objective expected utility theory defended it against the Allais Paradox by shifting their emphasis from the alleged descriptive power of the theory to its alleged normative power.

The same points can be made about the particular component of subjective expected utility theory that forms the central topic of this chapter—namely the notion of states of nature. It is one thing to assert that the states of nature approach offers a useful normative framework for decision making. It is quite a different thing to assert that, for the most part, this is how individuals actually do go about making decisions in the absence of probabilistic information. We shall consider each of these two domains in turn—in each case, with the goal of identifying the proper scientific criteria for states.

Criteria for Normative Applications

Savage's omelet example effectively shows how representing nature's underlying uncertainty by a set of “states,” then representing one's alternative courses of action as “acts” that map these states into their respective consequences, can serve to organize a decision problem and make it easier to see exactly how one's beliefs (the state likelihoods) and risk preferences should enter into the problem. For proper normative application, this first step—namely, the specification of the states—must satisfy three properties:

- 1) The alternative states must be mutually exclusive—that is, no two distinct states can simultaneously occur. Thus, it would not have been correct to list “egg is rotten” and “five-egg omelet” as two distinct states, since it is possible that these could simultaneously occur.

- 2) The family of states must be exhaustive—that is, whatever happens, at least one of the states can be said to have occurred. Although it is at the same logical level as the previous criterion (mutual exclusivity), the exhaustiveness criterion is much more difficult—and some would argue, actually impossible—to guarantee in practice. For example, if you cracked the sixth egg into the bowl and found that it was actually hollow, then neither of the two states in the Savage table could be said to have occurred, since neither of the first-row consequences would be realized (you would not have a six-egg omelet, nor would you have destroyed the other five eggs). When the decision maker has reason to “expect the unexpected,” the exhaustivity requirement cannot necessarily be achieved, and the best one can do is specify a final, catch-all state, with a label like “none of the above,” and a very ill-defined consequence.
- 3) The states must represent nature’s exogenous uncertainty, so their likelihoods cannot be affected by the individual’s choice of act. This issue can be illustrated by a simple example involving the decision whether or not to install a lightning rod on one’s house. Naturally, the relevant occurrences are the two mutually exclusive results {“house burns down,” “house doesn’t burn down”}. But since installing a lightning rod will clearly alter the respective likelihoods of these occurrences, can we really specify states of nature that are independent of the decision maker’s action? The answer is illustrated in the following table, which makes it clear that “house burns down” and “house doesn’t burn down” are not the states at all, but rather, part of the consequences, and clarifies that the effect of installing a lightning rod—as with any subjective act—is the outcome of an interaction between the act and an exogenous state of nature.

Act	State		
	Big lightning strike	Small lightning strike	No lightning strike
Lightning rod	House burns down, paid for rod	House doesn’t burn, paid for rod	House doesn’t burn, paid for rod
No lightning rod	House burns down, didn’t pay for rod	House burns down, didn’t pay for rod	House doesn’t burn, didn’t pay for rod

Do Decision Makers See the State Space or Do They Construct the State Space?

Is an individual who uses states of nature in normative decision making working with exogenous objects that they observe, or with endogenous objects that they construct? In one sense, this question is either subsidiary to, or equivalent to, the question of whether they are selecting from a menu of alternatives (subjective acts) that they observe as being available to them, or from a menu of alternatives that they have thought up or devised. Viewed in this larger sense, the question of whether the alternatives are observed or constructed is seen to have nothing to do with whether the choice happens to involve uncertainty at all, and indeed, the question may be equivalent to the classic question of whether Alexander Graham Bell discovered the idea for a telephone or invented this idea. In any case, I cannot derive any implications of this issue that pertain to the use of states of nature for normative purposes.

ISSUES OF OBSERVABILITY, CLASSIFICATION, AND MEASUREMENT

Independent Observability and the Exogeneity of States

Although the question of whether states are “exogenous and observed” versus “endogenous and constructed” does not seem to matter in a context of normative decision making, it matters a great deal for their relevance in descriptive science, for the types of reasons discussed in the “Scientific Considerations in the Theory of Choice” section. There we argued that economics had made greater scientific achievements than, say, anthropology because variables like prices and income were easier to measure than variables like trust or group identification. But if for some reason it should turn out that the full price of apples only exists in the eye of the consumer and is not independently observable, then this advantage is lost. This might be the case if the acquisition of a commodity involves a time cost, set-up cost, or transaction cost that is observable to the consumer, but not to the outside

observer. Note that an inability to observe the true price—an independent variable—poses scientific problems even if we can still observe the exact amount purchased—the dependent variable—since it impedes our ability to observe the relationship between the two (the true demand function).

In the context of choice under uncertainty, such a problem would arise whenever the state space used by the decision maker did not correspond to the state space hypothesized by the scientific observer. In some sense, this is less likely to happen if the states are exogenous objects that are observed than if they are endogenous objects that must be constructed. But even in the former case, there is the possibility that the decision maker observes either a finer or a coarser set of states than does the scientist. Ultimately, the question reduces the scientist's ability to view the set of actions available to the decision maker—the left-hand columns in the above decision tables—and correctly predict the decision maker's specification (be it an observation or construction) of both the upper row and cell entries. Where this can and cannot be done is an empirical question.

***Ex Ante* Observability versus *Ex Post* Observability of States**

Distinct from the question of whether the scientist can observe the set of states used by the decision maker is the question of whether the scientist can observe the realized state, or exactly when the scientist can observe the realized state. For example, in the case of choice under certainty—say, the demand for apples—it is clearly more important to be able to observe the price of apples that the consumer actually faces upon arriving at the supermarket, than to know the consumer's prior expectations of what this price might be. But interestingly enough, for choice under uncertainty, the ability to observe the state space before the fact is of much greater importance than the ability to observe the realized state. The reason is that choice under uncertainty is by definition *ex ante*, and only depends upon *ex ante* features of the decision problem, namely the state space and the set of available subjective acts over this space. A scientist who correctly gleans the decision maker's formulation of these concepts, who knows his beliefs over the likelihoods of the states, and who knows his attitudes toward risk, will be able to correctly predict his decision—a decision that by definition

must be made before, and hence cannot be influenced by, the actual realization of the state. In both the omelet and the lightning rod examples, *ex post* knowledge of the realized state is of no further predictive use for the scientist, except for possible future decisions, via its effect on the specification of a state space for some subsequent decision, and/or likelihood beliefs over this space.

Issues of Classification and Measurement

Are qualitative state spaces likely to be more or less subject to the above types of observability issues than quantitative state spaces? As the example of the unobserved apple price illustrates, even real-valued independent variables—real-valued commodity prices or real-valued states of nature—are subject to these issues in principle. On the other hand, decisions where the state space is more naturally quantitative are probably less subject to these specification difficulties than decisions where the state space is more naturally qualitative. For example, compare the uncertainty related to investing in a domestic farming company compared to the uncertainty related to investing in a similar company located in a politically unstable foreign country. In the former case, the state space probably only has few dimensions, all of which are quantitative: the average temperature over the growing season, the average rainfall over the season, and average output prices at harvest. In the latter case, the most significant sources of uncertainty may be subjective—the particular political party that comes to power and its subsequent choice of expropriation policy. There is every reason to think that the scientist will do a much better job of modeling the decision maker's problem formulation in the first case than in the second case. Indeed, in the following section we shall see that measurable, as opposed to qualitative, state spaces can actually serve to bring some mathematical structure of objective uncertainty into a purely subjective setting.

ISSUES OF STRUCTURE: ALMOST-OBJECTIVE UNCERTAINTY

As noted, an important feature of objective uncertainty is that it allows us to apply the analytical tools of probability theory, such as combinatorics, the Central Limit Theorem, the Law of Large Numbers, and Chebyshev's Inequality. Furthermore, since objective probabilities are part of the objects of choice themselves, these types of results can be invoked independently of, and prior to, any knowledge of the individual's attitude toward risk. Thus, for example, the conditions under which the sum of two independent objective lotteries will have the same distribution as (and thus presumably be indifferent to) some third lottery will be the same for all individuals. Such results have the same character as arbitrage results in portfolio theory, which hold independently of risk preferences and hence yield extremely powerful results.

But in some sense, this strength of the objective framework is also its greatest weakness: it imposes too much uniformity of beliefs across individuals, and in many cases, too much structure on each individual's own beliefs. In contrast with preferences over objective lotteries, preferences over real-world subjective prospects are subject to the following three phenomena:

- 1) Individuals may have different subjective likelihoods for the same event (diverse beliefs).
- 2) Individuals' beliefs may not be representable by probabilities at all, with some (or all) events being considered ambiguous (absence of probabilistic sophistication).
- 3) Individuals' outcome preferences may depend upon the source of uncertainty itself (outcome preferences may be state-dependent).

Nevertheless, it turns out that if the state space has a Euclidean structure and preferences are smooth in the events in the sense described in Appendix 2A, then features of "objective" uncertainty will emerge even in a purely subjective setting. In Appendix 2B, we sketch out the intuition of these results—readers wishing a formal development are referred to Machina (2001).

We can summarize the scientific implications of almost-objective uncertainty as follows: In the more traditional approach to uncertainty (e.g., Anscombe and Aumann 1963), the world presented two qualitatively different types of uncertainty: uncertain processes (such as perfectly balanced roulette wheels) that only generated idealized, purely *objective* events for which all agents held common beliefs and probabilistically sophisticated betting preferences; and uncertain processes (such as tomorrow's temperature or rainfall level) that only generated purely *subjective* events, where individuals typically differed in their likelihood beliefs, or had no likelihood beliefs, and could be state-dependent. On the other hand, according to the concepts presented in Appendix 2B, once a purely subjective state space is given a Euclidean structure and preferences are assumed to be smooth in the events, there exist events that arbitrarily closely approximate all the properties of classical "objective events" for all decision makers, in spite of any interpersonal differences in beliefs, lack of probabilistic sophistication, or state-dependence. Furthermore, once standard "objective randomizing devices" are reexamined, they are seen to depend precisely on these type of "almost-objective events."

Given the traditional (e.g., Savage 1954) approach of positing an almost completely unrestricted subjective state space and no event-smoothness, the "Euclidean state space + event-smoothness" approach advocated in the previous paragraph might seem overly strong. But in fact, it is well within standard economic practice. Standard consumer theory under certainty requires no structure at all on a family of objects of choice in order to axiomatize an ordinal utility function over these objects. Debreu's (1954) original topological assumptions were later shown to be unnecessary by Kreps (1988, pp. 25–26). But the workhorse concepts of competitive prices, marginal rates of substitution, demand functions, and the Slutsky equation do not emerge until we assume a Euclidean structure for these objects (vector "commodity bundles" and a Euclidean "commodity space") and/or smooth preferences over this space. Under uncertainty, restricting ourselves to a Euclidean state space amounts to nothing more than restricting ourselves to subjective uncertainty that appears in the form of random variables (such as temperature or random prices). And for the types of reasons discussed earlier in this chapter, real- or vector-valued states of nature are much more likely to be commonly observable and com-

only measurable than are states of nature that are elements of some more abstract space.

Just as science in general has progressed most rapidly when it has been able to quantify and measure the natural world, research in uncertain preferences and beliefs will further progress most rapidly to the extent we are able to quantify and measure the objects we call “states of nature” or “states of the world.”

Notes

I would like to thank Ted Groves, Donald Meyer, and Joel Sobel for helpful comments. This material is based upon work supported by the National Science Foundation under Grant No. 9870894. All opinions and errors are my own.

1. If the outcomes x describe monetary payoffs, then the standard monotonicity assumptions are that $V(x_1, p_1; \dots; x_m, p_m)$ is increasing in each of the variables x_1, \dots, x_m , and also increasing whenever p_i is increased at the expense of p_j for some pair of outcomes $x_i > x_j$ (that is, whenever probability mass is shifted from a lower to a higher outcome).
2. For example, an expected utility preference function $U(x_i) \cdot p_1 + \dots + U(x_m) \cdot p_m$ will be risk averse if and only if its *coefficient* with respect to $\text{prob}(x)$ (that is, the value $U(x)$) is a concave function of wealth. Correspondingly, a non-expected utility preference function $V(\mathbf{P}) = V(x_1, p_1; \dots; x_m, p_m)$ will be risk averse if and only if its *partial derivative* with respect to $\text{prob}(x)$ (that is, the value $\partial V(\mathbf{P})/\partial \text{prob}(x)$) is a concave function of wealth.
3. Such as the *Choquet expected utility* form

$$W_{\text{Choquet}}(x_1 \text{ on } E_1; \dots; x_m \text{ on } E_m) \equiv \sum_{i=1}^m U(x_i) \cdot [C(E_1 \cup \dots \cup E_i) - C(E_1 \cup \dots \cup E_{i-1})]$$

for some utility function $U(\cdot)$ and *capacity* (monotonic non-additive measure) $C(\cdot)$, where the outcomes are labeled so that $x_1 < \dots < x_m$ (e.g., Gilboa 1987; Schmeidler 1989; Wakker 1989, 1990; Gilboa and Schmeidler 1994), or the *maxmin expected utility* form

$$W_{\text{maxmin}}(x_1 \text{ on } E_1; \dots; x_m \text{ on } E_m) \equiv \min_{\tau \in T} \int_S U(f(s)) \cdot d\mu_\tau(s) \equiv \min_{\tau \in T} \sum_{i=1}^m U(x_i) \cdot \mu_\tau(E_i)$$

for some utility function $U(\cdot)$ and *family* $\{\mu_\tau(\cdot) \mid \tau \in T\}$ of probability measures on S (e.g., Gärdenfors and Sahlin 1982, 1983; Cohen and Jaffray 1985; Gilboa and Schmeidler 1989).

Appendix 2A

Properties of the Smooth Function Approach to Non-probabilistically Sophisticated Models

This approach starts by equivalently reexpressing each act $f(\cdot) = [x_1 \text{ on } E_1; \dots; x_n \text{ on } E_n]$ in the form $f(\cdot) = [\dots; x \text{ on } f^{-1}(x); \dots] = [\dots; x \text{ on } E_x; \dots]$, as x ranges over all possible outcomes $x \in \mathcal{X}$. The preference functions $W_{SEU}(\cdot)$ and $W_{SDEU}(\cdot)$ can then be expressed in the event-additive forms

$$W_{SEU}(\dots; x \text{ on } E_x; \dots) \equiv \sum_{x \in \mathcal{X}} \chi \Phi_x(E_x) \quad \text{where} \quad \Phi_x(E) \stackrel{\text{def}}{=} U(x) \cdot \mu(E)$$

$$W_{SDEU}(\dots; x \text{ on } E_x; \dots) \equiv \sum_{x \in \mathcal{X}} \chi \Phi_x(E_x) \quad \text{where} \quad \Phi_x(E) \stackrel{\text{def}}{=} \int_E U(x|s) \cdot d\mu(s)$$

where the event E_x attached to each outcome x is evaluated by an additive evaluation measure $\Phi_x(\cdot)$, which is the subjective analogue of objective expected utility's probability coefficient $U(x)$.

Just as linearity in a set of variables implies linearity in their changes, event-additive functions like $W_{SEU}(\dots; x \text{ on } E_x; \dots) = \sum_{x \in \mathcal{X}} \chi \Phi_x(E_x)$ and $W_{SDEU}(\dots; x \text{ on } E_x; \dots) = \sum_{x \in \mathcal{X}} \chi \Phi_x(E_x)$ will also be additive in event changes ("growth and shrinkage sets"). That is, their ranking of two acts $f(\cdot) = [\dots; x \text{ on } E_x; \dots]$ versus $f^*(\cdot) = [\dots; x \text{ on } E_x^*; \dots]$ is determined by the additive formulas

$$W_{SEU}(f^*(\cdot)) - W_{SEU}(f(\cdot)) \equiv \sum_{x \in \mathcal{X}} \Phi_x(E_x^*) - \sum_{x \in \mathcal{X}} \Phi_x(E_x) \equiv \sum_{x \in \mathcal{X}} \Phi_x(E_x^* - E_x) - \sum_{x \in \mathcal{X}} \Phi_x(E_x - E_x^*)$$

$$W_{SDEU}(f^*(\cdot)) - W_{SDEU}(f(\cdot)) \equiv \sum_{x \in \mathcal{X}} \Phi_x(E_x^*) - \sum_{x \in \mathcal{X}} \Phi_x(E_x) \equiv \sum_{x \in \mathcal{X}} \Phi_x(E_x^* - E_x) - \sum_{x \in \mathcal{X}} \Phi_x(E_x - E_x^*)$$

where for each x , its *growth set* in going from $f(\cdot)$ to $f^*(\cdot)$, namely the set $E_x^* - E_x$, is evaluated *positively* by x 's evaluation measure $\Phi_x(\cdot)$, and its *shrinkage set*, namely the set $E_x - E_x^*$, is evaluated *negatively* by $\Phi_x(\cdot)$.

Just as differentiability in objective probabilities can be defined as *local linearity in probability changes*, smoothness in subjective events can be defined as *local additivity in event changes*. That is, one can define a general preference function $W(\dots; x \text{ on } E_x; \dots)$ to be event-differentiable if at each act $f(\cdot)$ it possesses a family of local evaluation measures $\{\Phi_x(\cdot; f) \mid x \in \mathcal{X}\}$ such that

$W(\cdot)$ evaluates small event changes from $f(\cdot)$ in the following locally additive manner:

$$W(f^*(\cdot)) - W(f(\cdot)) \equiv \sum_{x \in \mathcal{X}} \Phi_x(E_x^* - E_x; f) - \sum_{x \in \mathcal{X}} \Phi_x(E_x - E_x^*; f) + o(\delta(f^*(\cdot), f(\cdot)))$$

where the distance function $\delta(f^*(\cdot), f(\cdot))$ between acts has the property that it shrinks to zero as the change sets $E_x^* - E_x$ and $E_x - E_x^*$ all shrink to zero, and (as with any definition of differentiability) $o(\cdot)$ denotes a function that is of higher order than its argument. In Machina (2002), I have shown how this calculus of events can be applied to establish the robustness of most of classical state-independent and state-dependent subjective expected utility theory and subjective probability theory to general event-smooth (but not necessarily either expected utility or probabilistically sophisticated) preference functions $W(\cdot)$ over subjective acts.

Appendix 2B

Almost-Objective Uncertainty

PROPERTIES OF PURELY OBJECTIVE EVENTS

We begin by contrasting the three properties of subjective events—diverse interpersonal beliefs, possible absence of probabilistic sophistication, and possible state-dependence—with the following four characteristic properties of idealized, exogenous “purely objective” events:

- *Unanimous, outcome-invariant revealed likelihoods*: In contrast with the above-listed properties of subjective events, all individuals exhibit identical, outcome-invariant revealed likelihoods over purely objective events—corresponding to their objective probabilities.
- *Independence from subjective realizations*: In the presence of joint objective \times subjective uncertainty, purely objective events are independent of the realization of subjective events. Thus, the events generated by an exogenous objective coin, die, or roulette wheel are invariant to whether any given subjective event E does or does not occur.
- *Probabilistic sophistication over objective lotteries*: It is almost a truism that all individuals evaluate objective lotteries $\mathbf{P} = (x_1, p_1; \dots; x_m, p_m)$ solely according to their outcomes and corresponding objective likelihoods, via some preference function $V(x_1, p_1; \dots; x_m, p_m)$.
- *Reduction of objective \times subjective uncertainty*: Standard reduction of compound uncertainty assumptions imply that individuals evaluate any objective mixture of subjective acts $\alpha \cdot f(\cdot) + (1-\alpha) \cdot f^*(\cdot) = \alpha \cdot [x_1 \text{ on } E_1; \dots; x_m \text{ on } E_m] + (1-\alpha) \cdot [x_1^* \text{ on } E_1^*; \dots; x_m^* \text{ on } E_m^*]$ solely according to its induced map $[\dots; (x_i, \alpha; x_j^*, 1-\alpha) \text{ on } E_i \cap E_j^*; \dots]$ from events to lotteries.

The above features of objective uncertainty apply to all individuals, whether or not they are expected utility, state-independent or probabilistically sophisticated. The following properties additionally hold for probabilistically sophisticated individuals and expected utility individuals:

- *Under probabilistic sophistication, independence of objective and subjective likelihoods:* If the individual is probabilistically sophisticated with probability measure $\mu(\cdot)$ over subjective events, these likelihoods are independent of exogenous objective events, and vice versa.
- *Under expected utility, linearity in objective likelihoods:* Expected utility is linear in objective probabilities ($V_{EU}(x_1, p_1; \dots; x_m, p_m) \equiv \sum_{i=1}^m U(x_i) \cdot p_i$) and in objective mixtures of lotteries ($V_{EU}(\alpha \cdot \mathbf{P} + (1-\alpha) \cdot \mathbf{P}^*) \equiv \alpha \cdot V_{EU}(\mathbf{P}) + (1-\alpha) \cdot V_{EU}(\mathbf{P}^*)$). Under objective \times subjective uncertainty, expected utility is linear in objective mixtures of subjective acts: $W_{SEU}(\alpha \cdot f(\cdot) + (1-\alpha) \cdot f^*(\cdot)) \equiv \alpha \cdot W_{SEU}(f(\cdot)) + (1-\alpha) \cdot W_{SEU}(f^*(\cdot))$, and similarly for $W_{SDEU}(\cdot)$.

ALMOST-EQUALLY-LIKELY EVENTS AND ALMOST-FAIR BETS


As the above bullet lists indicate, the properties of purely subjective and purely objective events lie in stark contrast. Nevertheless, in a Euclidean state space $\mathcal{S} = [\underline{s}, \bar{s}] \subseteq \mathbb{R}^1$, some subjective events are closer to being objective than others. We illustrate this by an example which approximates what is surely the “canonical” objective event: namely, the flip of an exogenous, fair coin.

Denoting the events implied by this coin by with standard notation $\{H, T\}$, their characteristic property is that, for any pair of prizes $x^* > x$, all individuals will be indifferent between the bets $[x^* \text{ on } H; x \text{ on } T]$ and $[x \text{ on } H; x^* \text{ on } T]$. In contrast, for any subjective event E , ranking of the bets $[x^* \text{ on } E; x \text{ on } \sim E]$ versus $[x \text{ on } E; x^* \text{ on } \sim E]$ can differ across individuals (due to diverse beliefs), or can reverse if the prizes $x^* > x$ are replaced by $y^* > y$ (due to state-dependence).

However, consider the event E_n obtained by dividing the state space $\mathcal{S} = [\underline{s}, \bar{s}]$ into n equal-length intervals, and defining E_n as the union of the *odd*-numbered intervals (the complementary event $\sim E_n$ thus being the union of the even-numbered intervals). As the following diagram indicates, *regardless* of an individual’s particular subjective probability measure $\mu(\cdot)$ over the state space \mathcal{S} (indicated by its density function $m(\cdot)$ in the Figure 2.B1), as n approaches infinity, the individual will assign equal subjective probabilities of $1/2$ to each of the events E_n and $\sim E_n$, and hence be virtually indifferent between the bets $[x^* \text{ on } E_n; x \text{ on } \sim E_n]$ versus $[x \text{ on } E_n; x^* \text{ on } \sim E_n]$. State-dependent individuals will be similarly indifferent, and as shown in Machina (2001), as $n \rightarrow \infty$, all event-smooth individuals—whether or not they are expected utility, state independent, or even probabilistically sophisticated—will “reveal” E_n and $\sim E_n$ to be equally likely, via their indifference between any two bets of the form $[x^* \text{ on } E_n; x \text{ on } \sim E_n]$ versus $[x \text{ on } E_n; x^* \text{ on } \sim E_n]$. In other words, as $n \rightarrow \infty$,

the purely subjective events E_n and $\sim E_n$ —both subsets of the purely subjective state space S —take on the properties of exogenous objective 50:50 events.

Figure 2B.1 Example of a Subjective Probability Density Function



ALMOST-OBJECTIVE EVENTS, ACTS AND MIXTURES

It is clear that by dividing the state space $S = [\underline{s}, \bar{s}]$ into a large number of equal-length intervals and taking the union of every third interval, we could create an subjective event that approximates the properties of an exogenous event of probability 1/3, etc. We can extend and formalize this idea as follows: Given any sufficiently regular (e.g., finite interval union) subset \wp of the unit interval $[0,1]$ and any large n , partition S into n equal-length intervals $[0, 1/n)$, $[1/n, 2/n) \dots [(n-2)/n, (n-1)/n)$, $[(n-1)/n, 1]$, and define the almost-objective event $\wp \times_n S \subseteq S$ by

$$\wp \times_n S = \bigcup_{i=0}^{n-1} \left\{ \underline{s} + (i+\omega) \cdot \frac{\bar{s}-\underline{s}}{n} \mid \omega \in \wp \right\}$$

that is, as the union of \wp 's linear images into each of S 's n equal-length intervals. Thus, the event E_n illustrated in the previous figure is simply almost-objective event $[0, 1/2] \times_n S$.

By taking a partition $\{\wp_1, \dots, \wp_m\}$ of the unit interval we can create almost-objective partitions $\{\wp_1 \times_n S, \dots, \wp_m \times_n S\}$ of the state space S , and in turn define almost-objective acts $[x_1 \text{ on } \wp_1 \times_n S; \dots ; x_m \text{ on } \wp_m \times_n S]$. The almost-fair bets of the previous subsection are seen to be the almost-objective acts $[x^* \text{ on } [0, 1/2] \times_n S; x \text{ on } (1/2, 1] \times_n S]$ and $[x \text{ on } [0, 1/2] \times_n S; x^* \text{ on } (1/2, 1] \times_n S]$. Finally, given two subjective acts $f(\cdot) = [x_1 \text{ on } E_1; \dots ; x_m \text{ on } E_m]$ and $f^*(\cdot) =$

$[x_1^*$ on E_1^* ; ... ; x_m^* on E_m^*] and any $\wp, \sim\wp \subseteq [0,1]$, we can define the al-

most-objective mixture $[f(\cdot)$ on $\wp \times_n \mathcal{S}$; $f^*(\cdot)$ on $\sim\wp \times_n \mathcal{S}]$ of $f(\cdot)$ and $f^*(\cdot)$.

BELIEFS AND BETTING PREFERENCES OVER ALMOST-OBJECTIVE EVENTS

As with the almost-equally likely events defined above, as $n \rightarrow \infty$ all event-smooth individuals will exhibit identical revealed likelihood beliefs over any almost-objective event $\wp \times_n \mathcal{S}$ essentially treating it as an exogenous objective event, with a probability given by the total length $\lambda(\wp)$ of the subset $\wp \subseteq [0, 1]$. That is to say, given any event-smooth preference function $W(\cdot)$ over subjective acts—whether or not it is expected utility/non-expected utility, state-independent/state-dependent, or probabilistically sophisticated/non-probabilistically sophisticated—outcomes $x^* > x$, disjoint subsets $\wp, \hat{\wp} \subseteq [0,1]$ with $\lambda(\wp) > \lambda(\hat{\wp})$, and subjective act $f(\cdot)$, $W(\cdot)$ will exhibit

$$\lim_{n \rightarrow \infty} W(x^* \text{ on } \wp \times_n \mathcal{S}; x \text{ on } \hat{\wp} \times_n \mathcal{S}; f(s) \text{ elsewhere}) >$$

$$\lim_{n \rightarrow \infty} W(x \text{ on } \wp \times_n \mathcal{S}; x^* \text{ on } \hat{\wp} \times_n \mathcal{S}; f(s) \text{ elsewhere});$$

that is, holding the payoffs elsewhere constant, all event-smooth individuals are unanimous in their preference for staking the greater of two prizes on the event $\wp \times_n \mathcal{S}$ and the lesser on $\hat{\wp} \times_n \mathcal{S}$, rather than the other way around. Thus, while we have seen that typical subjective events need not have probabilities at all, much less unanimously agreed-upon probabilities, as $n \rightarrow \infty$ there *will* be such unanimous agreement on the comparative likelihoods of $\wp \times_n$ versus $\hat{\wp} \times_n \mathcal{S}$.

The idea that some subjective events come close to exhibiting objective properties is not new, and precursors of the almost-equal-likelihood example date back at least to Poincaré (1912). Nor are almost-objective events merely a technical curiosum—in fact, most real-world “objective randomization devices” are actually examples of the use of almost-objective events to convert non-probabilistic subjective uncertainty to (almost-) objective uncertainty. To see this, consider the simple example of a game show spinner divided into a large number of alternating red and black sectors of equal angular size. Is it correct to say that the spin of such a wheel is an “objective process”? If so, then it would follow that all individuals would have the same beliefs over all events

defined over this process. But how much agreement will there be on the likelihood of the event that “the wheel spins more than 20 revolutions before finally stopping”?

Viewed from this perspective, the behavior of the wheel—its exact number of revolutions and therefore the color of the sector that finishes opposite the pointer—is a subjective process, where the state of nature is the amount of force applied to the spin. Individuals will surely disagree on their subjective probabilities of an event like “the force will be enough to generate at least 20 revolutions,” and some may not be able to attach any subjective probability at all to this event. But if we plot the state (the initial force of the spin) on the horizontal axis of the previous diagram, then an event such as “the force will lead the wheel to stop with the pointer opposite a black sector” is seen to be an almost-objective event of the type illustrated in the figure, which is why even individuals who disagree on the likelihood of “more than 20 spins” will nevertheless agree on the likelihood of “black.” In other words, it is not the *process* of spinning the wheel that is either “subjective” or “objective,” but rather the different *events* defined on this process that are either subjective or (almost-) objective. A little thought will reveal that virtually all standard physical randomization devices used to generate “objective” likelihoods share this property of being based on a subjectively uncertain (and hence non-probabilistic) state variable (or variables), but working with periodic, “almost-objective” events defined over the state variable.

The above argument shows that with a structured (essentially Euclidean) state space and the property of event-smooth preferences, there exists a substratum of events that arbitrarily closely approximate the first of the four above-listed properties of purely objective events, namely the property of unanimous, outcome-invariant revealed likelihoods. In Machina (2001) I have shown that such events, and the acts and mixtures based on them, also arbitrarily closely approximate the other three listed properties of idealized “purely objective” events. That is, as $n \rightarrow \infty$:

- Each individual (probabilistically sophisticated or otherwise) will view all almost-objective events as independent of each purely subjective event, in the sense that for all disjoint $\wp, \hat{\wp} \subseteq [0,1]$ and each $E \subseteq \mathcal{S}$, they will have the same revealed likelihood rankings (i.e., betting preferences) over the joint events $(\wp \times_n \mathcal{S}) \cap E$ versus $(\hat{\wp} \times_n \mathcal{S}) \cap E$ as they do over the events $\wp \times_n \mathcal{S}$ versus $\hat{\wp} \times_n \mathcal{S}$ (in each case, corresponding to the relative values of $\lambda(\wp)$ versus $\lambda(\hat{\wp})$).

- Although individuals needn't be probabilistically sophisticated over subjective acts in general, they will be probabilistically sophisticated over almost-objective acts. That is, each $W(\cdot)$ will have a corresponding preference function $V_W(\cdot)$ over lotteries such that

$$\lim_{n \rightarrow \infty} W \left(x_1 \text{ on } \wp_1 \times_n S; \dots; x_m \text{ on } \wp_m \times_n S \right) = \lim_{n \rightarrow \infty} V_W \left(x_1, \lambda(\wp_1); \dots; x_m, \lambda(\wp_m) \right).$$

- Each individual (probabilistically sophisticated or otherwise) satisfies the reduction of compound uncertainty property for almost-objective mixtures of acts. Thus if $\left[f_1(\cdot) \text{ on } \wp_1 \times_n S; \dots; f_m(\cdot) \text{ on } \wp_m \times_n S \right]$ and $\left[\hat{f}_1(\cdot) \text{ on } \hat{\wp}_1 \times_n S; \dots; \hat{f}_m(\cdot) \text{ on } \hat{\wp}_m \times_n S \right]$ induce almost-objectively equivalent subacts over each event in the common refinement of

$f_1(\cdot), \dots, f_m(\cdot), \hat{f}_1(\cdot), \dots, \hat{f}_m(\cdot)$, then

$$\begin{aligned} \lim_{n \rightarrow \infty} W \left(f_1(\cdot) \text{ on } \wp_1 \times_n S; \dots; f_m(\cdot) \text{ on } \wp_m \times_n S \right) = \\ \lim_{n \rightarrow \infty} W \left(\hat{f}_1(\cdot) \text{ on } \hat{\wp}_1 \times_n S; \dots; \hat{f}_m(\cdot) \text{ on } \hat{\wp}_m \times_n S \right). \end{aligned}$$

The following properties of almost-objective uncertainty additionally hold for probabilistically sophisticated individuals and expected utility individuals:

- Each probabilistically sophisticated individual with subjective probability measure $\mu(\cdot)$ will view *all purely subjective events as independent of each almost-objective event*, in the sense that for all $E, \hat{E} \subseteq S$ and each $\wp \subseteq [0, 1]$, they will have the same revealed likelihood rankings (betting preferences) over the joint events $(\wp \times_n S) \cap E$ versus $(\wp \times_n S) \cap \hat{E}$ as they do over the events E versus \hat{E} (in each case, corresponding to the relative values of $\mu(E)$ versus $\mu(\hat{E})$).
- Each expected utility maximizer will be *linear in almost-objective probabilities and almost-objective mixtures of subjective lotteries*, i.e.,

$$\lim_{n \rightarrow \infty} W_{SEU} \left(x_1 \text{ on } \wp_1 \times_n S; \dots; x_m \text{ on } \wp_m \times_n S \right) = \sum_{i=1}^m \lambda(\wp_i) \cdot W_{SEU} \left(x_i \text{ on } S \right)$$

$$\lim_{n \rightarrow \infty} W_{SEU} \left(f_1(\cdot) \text{ on } \wp_1 \times_n S; \dots; f_m(\cdot) \text{ on } \wp_m \times_n S \right) = \sum_{i=1}^m \lambda(\wp_i) \cdot W_{SEU} \left(f_i(\cdot) \right)$$

and similarly for $W_{SDEU}(\cdot)$.

References

- Allais, M. 1953. "Le Comportement de l'Homme Rationnel devant le Risque, Critique des Postulats et Axiomes de l'Ecole Américaine." *Econometrica* 21(4): 503–546.
- Anscombe F., and R. Aumann. 1963. "A Definition of Subjective Probability." *Annals of Mathematical Statistics* 34(1): 199–205.
- Bernoulli, D. [1738] 1954. "Specimen Theoriae Novae de Mensura Sortis," *Commentarii Academiae Scientiarum Imperialis Petropolitanae* [Papers of the Imperial Academy of Sciences in Petersburg] V, 175-192. English translation: "Exposition of a New Theory on the Measurement of Risk." *Econometrica* 22(1): 23–36.
- Billingsley, P. 1986. *Probability and Measure*. 2d ed. New York: John Wiley and Sons.
- Chew, S. 1983. "A Generalization of the Quasilinear Mean With Applications to the Measurement of Income Inequality and Decision Theory Resolving the Allais Paradox." *Econometrica* 51(4): 1065–1092.
- Cohen, M., and J. Jaffray. 1985. "Decision Making in the Case of Mixed Uncertainty: A Normative Model." *Journal of Mathematical Psychology* 29(4): 428–442.
- Debreu, G. 1954. "Representation of a Preference Ordering by a Numerical Function." In *Decision Processes*, R. Thrall, C. Boombs, and R. Davis, eds. New York: John Wiley and Sons, pp. 159–165.
- Edwards, W. 1955. "The Prediction of Decisions among Bets." *Journal of Experimental Psychology* 50(3): 201–214.
- . 1962. "Subjective Probabilities Inferred From Decisions." *Psychological Review* 69(2): 109–135.
- Ellsberg, D. 1961. "Risk, Ambiguity, and the Savage Axioms." *Quarterly Journal of Economics* 75: 643–669.
- . 2001. *Risk, Ambiguity and Decision*. New York: Garland Publishing, Inc.
- Feller, W. 1968. *An Introduction to Probability Theory and Its Applications, Volume I*. 3d Ed. New York: John Wiley and Sons.
- . 1971. *An Introduction to Probability Theory and Its Applications, Volume II*. 2d Ed. New York: John Wiley and Sons.
- Fishburn, P. 1983. "Transitive Measurable Utility." *Journal of Economic Theory* 31(2): 293–317.

- Gärdenfors, P., and N. Sahlin. 1982. "Unreliable Probabilities, Risk Taking, and Decision Making." *Synthese* 53(3): 361–386.
- . 1983. "Decision Making with Unreliable Probabilities." *British Journal of Mathematical and Statistical Psychology* 36(2): 240–251.
- Gilboa, I. 1987. "Expected Utility with Purely Subjective Non-Additive Probabilities." *Journal of Mathematical Economics* 16(1): 65–88.
- Gilboa, I. and D. Schmeidler. 1989. "Maximum Expected Utility With a Non-Unique Prior." *Journal of Mathematical Economics* 18(2): 141–153.
- . 1994. "Additive Representations of Non-Additive Measures and the Choquet Integral." *Annals of Operations Research* 52: 43–65.
- Kahneman, D., and A. Tversky. 1979. "Prospect Theory: An Analysis of Decision Under Risk." *Econometrica* 47(2): 263–291.
- Kreps, D. 1988. *Notes on the Theory of Choice*. Boulder, Colorado: Westview Press.
- . 1990. *A Course in Microeconomic Theory*. Princeton: Princeton University Press.
- Machina, M. 1982. "'Expected Utility' Analysis Without the Independence Axiom." *Econometrica* 50(2): 277–323.
- . 1983. "Generalized Expected Utility Analysis and the Nature of Observed Violations of the Independence Axiom." In *Foundations of Utility and Risk Theory with Applications*, B. Stigum and F. Wenstop, eds. Dordrecht: D. Reidel Publishing Co., pp. 263–293.
- . 1987. "Choice Under Uncertainty: Problems Solved and Unsolved." *Journal of Economic Perspectives* 1(1): 121–154.
- . 2001. "Almost-Objective Uncertainty." Working paper 2001-12. San Diego: University of California, Department of Economics.
- . 2002. "Robustifying the Classical Model of Risk Preferences and Beliefs." Working paper 2002-06. San Diego: University of California, Department of Economics.
- MacCrimmon, K., and S. Larsson. 1979. "Utility Theory: Axioms Versus 'Paradoxes'." In *Expected Utility Hypotheses and the Allais Paradox*, Allais and Hagen, eds. Dordrecht: D. Reidel Publishing Co., pp. 333–409.
- Mas-Colell, A., M. Whinston and J. Green. 1995. *Microeconomic Theory*. Oxford: Oxford University Press.
- Poincaré, H. 1912. *Calcul des Probabilités*. 2d ed. Paris: Gauthiers-Villars.
- Quiggin, J. 1982. "A Theory of Anticipated Utility." *Journal of Economic Behavior and Organization* 3(4): 323–343.
- Savage, L. 1954. *The Foundations of Statistics*. New York: John Wiley and Sons.
- Schmeidler, D. 1989. "Subjective Probability and Expected Utility without Additivity." *Econometrica* 57(3): 571–587.

- Varian, H. 1992. *Microeconomic Analysis*. 3d ed. New York: W.W. Norton & Co.
- Wakker, P. 1989. "Continuous Subjective Expected Utility With Non-Additive Probabilities." *Journal of Mathematical Economics* 18(1): 1–27.
- . 1990. "Under Stochastic Dominance Choquet-Expected Utility and Anticipated Utility are Identical." *Theory and Decision* 29(2): 119–132.
- Yaari, M. 1987. "The Dual Theory Of Choice Under Risk." *Econometrica* 55(1): 95–115.