BAYESIANISM WITHOUT PRIORS, ACTS WITHOUT CONSEQUENCES

Robert Nau

Fuqua School of Business, Duke University

www.duke.edu/~rnau

September 13, 2005

Abstract: A generalization of subjective expected utility is presented in which the primitives are a finite set of states of the world, a finite set of strategies available to the decision maker, and allocations of money. The model does not require explicit definitions of consequences ("states of the person"), nor does it rely on counterfactual preferences, nor does it emphasize the unique separation of prior probabilities from possibly-state-dependent utilities. Rather, preferences have an additively separable representation in which the valuation of outcomes of a decision or game is implicit in the state- and strategy-dependence of utility for money. This model provides an axiomatic foundation for Bayesian decision analysis and game theory in the tradition of de Finetti and Arrow-Debreu rather than Savage. The observable parameters of beliefs are risk neutral probabilities (betting rates for money) and in situations where the decision maker has no intrinsic interest or influence over an experiment given the truth or falsity of the hypothesis, her risk neutral probabilities and preferences among strategies are updated by application of Bayes' rule without the need to identify "true" prior probabilities.

Keywords. Subjective probability, state-dependent utility, state-preference theory, risk neutral probabilities

Acknowledgements: I am grateful to Edi Karni, Jim Smith, Peter Wakker, and participants at RUD 2005 for helpful comments.

1. INTRODUCTION

The theory of rational choice under uncertainty relies heavily on the subjective expected utility model axiomatized by Savage (1954) and Anscombe and Aumann (1963), which justifies the use of subjective probabilities and utilities to represent the beliefs and values of a rational decision maker. The SEU model is called "Bayesian" for its implications concerning the receipt of information: when the decision maker observes the result of an experiment, she updates her beliefs by multiplying her prior probability distribution by the likelihood function that summarizes the data, following Bayes' rule, and modifies her preferences accordingly.

The primitive elements of Savage's model are a set of "states of the world" representing uncertain future events and a set of "consequences" representing possible outcomes of a decision problem. In Savage's words, "a consequence is anything that may happen to the person" in terms of wealth, health, the enjoyment of life, or the fate of loved ones. An "act" (course of action) for the decision maker is then defined as a mapping of states to consequences or, in Anscombe and Aumann's variant, to objective lotteries over consequences. Thus, formally, the notion of a consequence precedes the notion of an act, even though ontologically it is the other way around: a *consequence* is literally what *follows from* the decision maker's choice of a feasible course of action in a possible state of the world.

In order to pin down subjective probabilities and utilities, Savage's theory requires the decision maker to contemplate a very large set of imaginary, counterfactual acts in which states are mapped to consequences in every possible way, such as an act in which she enjoys being alive in a state of the world in which she is supposed to be dead. Let $\boldsymbol{f}$ denote such an act, and let f_s denote the consequence that it yields in state s. Savage's axioms imply that preferences among acts are represented by a total utility function U of the form

$$U(\boldsymbol{f}) = \sum_{s \in S} p_s u(f_s)$$

where $\boldsymbol{p}$ is a unique subjective probability distribution over states and u is a state-independent utility function for consequences that is unique up to a positive affine transformation.

Several objections can be raised against the use of this model as a foundation for theories of economic behavior under uncertainty. First, the combinatorial complexity of the act space requires the preference axioms to bear enormous weight, particularly the axioms of completeness and independence which are controversial in their own right. Second, even if the model is accepted as an idealized representation of an individual's state of mind, the fact that its measurements refer to imaginary acts yielding privately-experienced consequences makes it epistemically implausible for different individuals to know each other's u or agree on $\boldsymbol{p}$, as is typically assumed in game theory. Third, even if everything else is granted, the state-independence of u is behaviorally untestable, which leaves the uniqueness of $\boldsymbol{p}$ in doubt. For, let $\boldsymbol{q}$ be any other probability distribution and let $v_s(x) = u(x)p_s/q_s$ be a state-dependent utility function for consequences, which differs from u only by a state-dependent scale factor. Then the decision maker's preferences are represented by the pair ($\boldsymbol{q}$, $\boldsymbol{v}$) as well as by the pair ($\boldsymbol{p}$, u), and there is no way to tell which is the "true" representation of beliefs and values.

The perverse logic of primitive consequences that can be attached to arbitrary states has been criticized by a number of authors, most notably by Aumann (1971), Shafer (1986), and Skiadas (1997ab). Aumann's questions provoked a famous response from Savage: "To some—perhaps to you—it will seem grotesque if I say that I should not mind being hung so long as it be done without damage to my health or reputation, but I think it desirable to adopt such language so that the danger of being hung can be contemplated in the framework of [*Foundations of*

Statistics].” Savage’s view has largely prevailed, so that it has become conventional in decision theory to define acts in terms of consequences rather than the other way around, and (notwithstanding the subversive possibility of state-dependent utility scale factors) this practice is often defended as being necessary for separating beliefs from values, insofar as consequences are supposedly the fundamental particles of value. The separation of beliefs from values is a prized modeling objective because it enables preferences to be parameterized in a parsimonious way and because it jibes with the distinction between the effects of *information* and *tastes* that is traditional in statistics and economics. For example, Bayesian statistical inference is mainly concerned with how subjective probabilities are modified by data (information), while game theory starts from a description of the “rules of the game” in terms of the utility functions (tastes) of the players, from which probabilities of outcomes are endogenously determined by solution concepts such as Nash equilibrium. The inability to uniquely separate probabilities from utilities would appear to be a serious problem for these theories. Schervish et al. (1990) write: “Much of probability theory and statistical theory deals solely with probabilities and not with utilities. If probabilities are only unique relative to a specified utility, then much of this theory is in doubt.” On the flip side of this coin, if utilities of outcomes are not determined by preferences unless probabilities are exogeneously specified, then the knowledge assumptions that underly noncooperative game theory are equally in doubt.

This paper presents an alternative model of choice under uncertainty whose primitives are states of the world, strategies (feasible courses of action) available to the decision maker, and allocations of money. To adopt the terminology of game theory, the events in question are “moves” by two players, one of whom is the decision maker and the other of whom is an adversary who might be “nature” or “the economy” or an intelligent competitor. (The generalization to more than two players is straightforward.) These events are assumed to be independently observable so as to provide a basis for monetary bets between the decision maker and outside observers, through which the parameters of her preferences can be made common knowledge and through which the rationality of her behavior ultimately can be judged. This choice of primitives circumvents the epistemic problems of Savage-type consequences, as well as related problems of approximating the "grand world" by "small worlds", and provides the foundation for a decision theory that is fully operational in the sense that it refers only to observable events and monetary transactions rather than imaginary private experiences.

Although money is the only reward that appears directly in the model, this does not mean that non-monetary consequences of states and strategies are ignored. Rather, the valuation of non-monetary (or otherwise unobservable) consequences is indirectly revealed through the decision maker's preferences among strategies paired with allocations of monetary wealth. Every combination of a state, a strategy, and a quantity of money is effectively treated as a distinct consequence, and there is no need for counterfactual acts. A cross-strategy form of the independence axiom ("strategic generalized triple cancellation") is used to obtain an additively separable utility representation in which prior beliefs are not uniquely separated from state-dependent utilities, and the observable parameters of the decision maker's beliefs are "risk neutral probabilities" instead of "true" subjective probabilities. However, these turn out to be sufficient for purposes of Bayesian statistical inference and the modeling of decisions, games, and markets along the lines of Nau and McCardle (1990, 1991) and Nau (1992, 1995ac, 2001, 2003). Indeed, this paper provides the minimal axiomatic foundation needed for the models developed in those earlier papers. It will be argued that *the core principle of Bayesianism is not the determination of a unique prior probability distribution, but rather the process of prior-to-*

posterior updating of an additively separable utility function through multiplication by a likelihood function that summarizes the data. In situations where the decision maker has no intrinsic interest or influence over an experiment given the truth or falsity of the hypothesis, her true likelihood function is observable by de Finetti's method, while her true prior probabilities need not be, and her preferences following the receipt of new information are modified by updating her risk neutral probabilities according to Bayes' rule.

2. THE MODEL

The modeling framework of the paper is essentially that of Arrow-Debreu state-preference theory, augmented by the introduction of a finite set of *strategies* (discrete courses of action or "moves") available to the decision maker. Let A denote the set of strategies (with generic elements a, b, ...) and let S denote a finite set of *states* (with generic elements r, s, ...). Non-empty subsets of S are *events*. Let $\boldsymbol{w}$ denote an *allocation*, i.e., a real-valued vector of observable state-dependent monetary payoffs to the decision maker, in addition to unobservable status quo wealth and any other personal consequences of strategies and states. The set of all possible allocations is a convex, rectangular subset of $\Re^{|S|}$ denoted by W and endowed with the $\Re^{|S|}$ topology. (Thus, for each state there is some interval of possible payoffs.) w_s denotes the payoff of $\boldsymbol{w}$ in state s and $\boldsymbol{w}_{-s}$ denotes the vector of payoffs in states other than s_s. Similarly, $\boldsymbol{w}_E$ denotes the vector of payoffs assigned by $\boldsymbol{w}$ to the states in event E, and $\boldsymbol{w}_{-E}$ denotes the payoffs in states in the complement of E. $\boldsymbol{w}_E\ \boldsymbol{w}^*_{-E}$ denotes an allocation that agrees with $\boldsymbol{w}$ in event E and agrees with $\boldsymbol{w}^*$ in all other states, whence by definition $\boldsymbol{w} = w_s\ \boldsymbol{w}_{-s} = \boldsymbol{w}_E\ \boldsymbol{w}_{-E}$.

A pair $(a, \boldsymbol{w})$ constitutes an *act* for the decision maker in which strategy $a \in A$ is implemented and allocation $\boldsymbol{w} \in W$ is received. Thus, for example, the decision maker might be offered various allocations in association with different strategies in order to explore her strengths of preference among strategies or to assess her risk attitudes. Or she might bet with other individuals or purchase contingent claims in order to hedge risks or finance a preferred strategy. Some acts may be only hypothetical (e.g., if they violate present budget constraints or markets are incomplete), but they are not logically impossible. An *outcome* is a triple (a, s, w), which is effectively treated as a unique consequence.

Assume that the decision maker has preferences over acts denoted by a binary relation $\succcurlyeq$. A strategy-dependent utility function $U_a(\boldsymbol{w})$ *represents* $\succcurlyeq$ if $(a, \boldsymbol{w}) \succcurlyeq (b, \boldsymbol{x}) \Leftrightarrow U_a(\boldsymbol{w}) \geq U_b(\boldsymbol{x})$. It will be shown that under suitable axioms $\succcurlyeq$ has the additive representation:

$$U_a(\boldsymbol{w}) = \sum_{s \in S} v_{as}(w_s) \tag{1}$$

where $\{v_{as}\}$ are strictly increasing cardinal *evaluation functions* for money that depend on both the state of nature and the strategy chosen by the decision maker. Valuations of non-monetary or otherwise unobservable consequences of acts are implicit in the state-and-strategy-dependence of the evaluation functions for money.[1] In particular, the intercepts $\{v_{as}(0)\}$ determine the decision maker's preferences among strategies in the absence of additional monetary compensation. Subjective prior probabilities do not appear in the model: there is no explicit separation of beliefs from values. Inasmuch as (1) is a generalization of subjective expected utility, it *could* be the case that the decision maker has prior beliefs represented by a subjective probability distribution on S that she uses to construct her preferences, in conjunction with a possibly-state-

[1] Skiadas (1997ab) introduces a somewhat similar framework of subjective uncertainty without primitive consequences. However, Skiadas' model differs from ours in that it does not give an explicit role to money, and his analysis is directed at other issues such as nonseparable preferences, disappointment and regret, and (in 1997b) the unique separation of probability from utility via an additional binary relation that compares events according to whether they yield "equally desirable" consequences under a given act, a type of comparison that does not correspond to a choice.

dependent utility function for money. That is, it could be the case that $v_{as}(w) = p_s u_{as}(w)$ where $\boldsymbol{p}$ is a subjective probability distribution that encodes prior beliefs and $\{u_{as}\}$ are Bernoulli utility functions that encode values or risk attitudes. It could even be the case that the subjective probabilities depend on the strategies, i.e., $v_{as}(w) = p_{as}u_{as}(w)$, as though the decision were able to influence the likelihoods of some events through her actions, in the spirit of Karni's (2005) model. However, the prior distribution need not be observable. Despite the general lack of determinacy of "true" subjective probabilities, the decision maker has unique, observable *risk neutral probabilities* that are determined by the derivatives of the state-and-strategy-dependent evaluation functions, whose central role in decision modeling will be discussed further in the sequel. The preference axioms and main representation theorem are as follows:

Axiom 1: (Ordering) $\succcurlyeq$ is a weak order (complete, transitive, reflexive), and for every strategy the induced preference relation over allocations is a continuous weak order.

Axiom 2: (Strict monotonicity) $w^* > w \Rightarrow (a, w^*_s\, \boldsymbol{w}_{-s}) \succ (a, \boldsymbol{w})$ for all $a, s, \boldsymbol{w}, w^*_s$.

Note that axiom 2 implies that no states are null, as well as that more money is strictly preferred to less.

Axiom 3: (Strategic generalized triple cancellation) For every pair of strategies a and b, allocations $\boldsymbol{w}$, $\boldsymbol{x}$, $\boldsymbol{y}$, $\boldsymbol{z}$, state s, and constants w^*, x^*: $(a, \boldsymbol{w}) \preccurlyeq (b, \boldsymbol{x})$ and $(a, w^*_s\, \boldsymbol{w}_{-s}) \succcurlyeq (b, x^*_s\, \boldsymbol{x}_{-s})$ and $(a, w_s\, \boldsymbol{y}_{-s}) \succcurlyeq (b, x_s\, \boldsymbol{z}_{-s}) \Rightarrow (a, w^*_s\, \boldsymbol{y}_{-s}) \succcurlyeq (b, x^*_s\, \boldsymbol{z}_{-s})$.

Axiom 3, which appears to be novel, is the key to the representation. In words, it requires that if changing w to w^* is no worse than changing x to x^* in state s when the background is a comparison of (a, $\boldsymbol{w}$) vs. (b, $\boldsymbol{x}$), then changing w to w^* cannot be strictly worse than changing x to x^* in state s when the background comparison is (a, w_s $\boldsymbol{y}_{-s}$) vs. (b, x_s $\boldsymbol{z}_{-s}$). In the special case where a=b, axiom 3 reduces to ordinary generalized triple cancellation (*c.f.* Wakker 1989), which implies the familiar coordinate independence axiom (Savage's P2) but is slightly stronger: it yields an additive representation of preferences for wealth under strategy a even in the case $|S|$=2 without the additional hexagon condition. The more general case where a≠b serves to link the representations for wealth preferences under a and wealth preferences under b.

> **Axiom 4**: (Overlap among strategies) There exist allocations $\{ \boldsymbol{w}^{(a)}, a \in A \}$ in the interior of W such that $(a, \boldsymbol{w}^{(a)}) \sim (b, \boldsymbol{w}^{(b)})$ for all strategies a and b.

The role of the this axiom (and its conditional extension 4* below) is to ensure that cardinal utility scales can be linked across strategies and that the "spectrum" of utility has no holes. It is somewhat stronger than is strictly necessary (a sequence of pairwise overlaps would suffice), but it is transparent and leads to simple proofs.

> **Theorem 1:** Axioms 1-4 hold iff ≽ is represented by a utility function U having the additive form (1) in which $\{v_{as}\}$ are strictly increasing, continuous evaluation functions for money, with $U_a(\boldsymbol{w}^{(a)}) = U_b(\boldsymbol{w}^{(b)})$ for all strategies a and b. U is unique up to positive affine transformations, and $\{v_{as}\}$ are unique up to joint transformations of the form $\alpha v_{as} + \beta_s + \gamma_{as}$ where α is an arbitrary positive scale factor and $\{\beta_s\}$ and $\{\gamma_{as}\}$ are arbitrary constants with $\{\gamma_{as}\}$ summing to zero for every a.

The arbitrary constant term $\beta_s + \gamma_{as}$ in the evaluation function may be interpreted as an unobservable *externality* imposed on the decision maker by action of the opponent (nature or whomever), consisting of a main effect β_s and a zero-sum interaction γ_{as}. (The interaction term can be constrained to some extent by conditional preferences, as discussed in the next section.) The fact that externalities are unobservable has profound implications for game theory, because externalities determine whether a solution that is individually rational is also socially rational, as in the prisoner's dilemma.

3. CONDITIONAL PREFERENCES

In Savage's theory, conditional preferences are induced by unconditional preferences on a rich set of (mostly imaginary) contingent acts. Let $\boldsymbol{f}$ and $\boldsymbol{g}$ be Savage acts and let E be any event whatsoever. If $\boldsymbol{f}_E\boldsymbol{h}_{-E} \succcurlyeq \boldsymbol{g}_E\boldsymbol{h}_{-E}$ for some other act $\boldsymbol{h}$, then by the independence axiom the same preference holds for all $\boldsymbol{h}$. (Here $\boldsymbol{f}_E\boldsymbol{h}_{-E}$ denotes the contingent act that agrees with $\boldsymbol{f}$ in all states in E and agrees with $\boldsymbol{h}$ elsewhere, so that by definition the two acts $\boldsymbol{f}_E\boldsymbol{h}_{-E}$ and $\boldsymbol{g}_E\boldsymbol{h}_{-E}$ agree with each other when E does not occur. Such acts can always be constructed in Savage's framework, because it includes all possible mappings of states to consequences.) $\boldsymbol{f}$ is then defined to be conditionally preferred to $\boldsymbol{g}$ given E, denoted $\boldsymbol{f} \succcurlyeq_E \boldsymbol{g}$, if $\boldsymbol{f}_E\boldsymbol{h}_{-E} \succcurlyeq \boldsymbol{g}_E\boldsymbol{h}_{-E}$ for any and all $\boldsymbol{h}$, which lays the groundwork for a corresponding definition of conditional probability. This conventional definition cannot be applied in the framework of §2 because, given an arbitrary event E and two acts $(a, \boldsymbol{w})$ and $(b, \boldsymbol{x})$ that involve distinct strategies, there is no way (yet) to construct two other related acts that agree with $(a, \boldsymbol{w})$ and $(b, \boldsymbol{x})$, respectively, if E occurs and agree *with each other* if E does not occur. However, an analogous definition can be used if some additional structure is imposed on the state space in order to construct a richer set of (feasible) contingent acts. Henceforth assume that the state space S can be partitioned as $S = H \times I = \{H_1, \ldots H_J\} \times \{I_1, \ldots, I_K\}$ where I is a set of "informational" events that are expected to be resolved before an act is chosen and H is a set of "hypotheses" or otherwise consequential events that are expected to be resolved only after an act has been chosen. This temporal structure requires that events in I be causally independent of the decision maker's action, while events in H need not be. In addition to the "pure" strategies in A, let the decision maker also be free to choose *contingent strategies* that are pegged to events in I. Let $\tilde{\boldsymbol{a}}: I \rightarrow A$ denote such a contingent strategy, and let $\tilde{a}(s)$ denote the pure strategy assigned by $\tilde{\boldsymbol{a}}$ to state s. Then the pair $(\tilde{\boldsymbol{a}}, \boldsymbol{w})$ is a *contingent act*, and the preference relation $\succcurlyeq$ and its axioms 1-3 are henceforth assumed to cover contingent as well as pure acts. If $\tilde{\boldsymbol{a}}$ and $\tilde{\boldsymbol{b}}$ are two contingent strategies and $E \subset I$, $\tilde{\boldsymbol{a}}_E\tilde{\boldsymbol{b}}_{-E}$ denotes the contingent strategy that agrees with $\tilde{\boldsymbol{a}}$ on E and agrees with $\tilde{\boldsymbol{b}}$ elsewhere. Preferences conditioned on informational events (only) are now defined and linked to unconditional preferences through:

Axiom 5: (Independence) For every event $E \subset I$, $(\tilde{\boldsymbol{a}}_E\tilde{\boldsymbol{c}}_{-E}, \boldsymbol{w}_E\boldsymbol{y}_{-E}) \succcurlyeq (\tilde{\boldsymbol{b}}_E\tilde{\boldsymbol{c}}_{-E}, \boldsymbol{x}_E\,\boldsymbol{y}_{-E}) \Leftrightarrow$ $(\tilde{\boldsymbol{a}}_E\tilde{\boldsymbol{d}}_{-E}, \boldsymbol{w}_{\boldsymbol{E}}\boldsymbol{z}_{-E}) \succcurlyeq (\tilde{\boldsymbol{b}}_E\tilde{\boldsymbol{d}}_{-E}, \boldsymbol{x}_E\boldsymbol{z}_{-E})$ for all $\tilde{\boldsymbol{a}}$, $\tilde{\boldsymbol{b}}$, $\tilde{\boldsymbol{c}}$, $\tilde{\boldsymbol{d}}$, $\boldsymbol{w}, \boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z}$

Definition: $(\tilde{\boldsymbol{a}}, \boldsymbol{w})$ is weakly *conditionally preferred* to $(\tilde{\boldsymbol{b}}, \boldsymbol{x})$ given E, denoted $(\tilde{\boldsymbol{a}}, \boldsymbol{w}) \succcurlyeq_E (\tilde{\boldsymbol{b}}, \boldsymbol{x})$, if $(\tilde{\boldsymbol{a}}_E\tilde{\boldsymbol{c}}_{-E}, \boldsymbol{w}_E\boldsymbol{y}_{-E}) \succcurlyeq (\tilde{\boldsymbol{b}}_E\tilde{\boldsymbol{c}}_{-E}, \boldsymbol{x}_{\boldsymbol{E}}\boldsymbol{y}_{-E})$ for some $\tilde{\boldsymbol{c}}$ and $\boldsymbol{y}$.

Axiom 6: (Aggregation) For disjoint events $E, F \subset I$: $(\tilde{\boldsymbol{a}}, \boldsymbol{w}) \succcurlyeq_E (\tilde{\boldsymbol{b}}, \boldsymbol{x})$ and $(\tilde{\boldsymbol{a}}, \boldsymbol{w}) \succcurlyeq_F (\tilde{\boldsymbol{b}}, \boldsymbol{x}) \Rightarrow (\tilde{\boldsymbol{a}}, \boldsymbol{w}) \succcurlyeq_{E \cup F} (\tilde{\boldsymbol{b}}, \boldsymbol{x})$

Axiom 5 rules out non-neutral attitudes toward disappointment, so that the evaluation of an act conditional on an event E is not affected by what the act would have yielded had E not occurred. Axiom 6 corresponds to the informal version of the sure thing principle sketched by Savage (1954, p. 21), which is referred to as "coherence" by Skiadas (1997ab). A different name is adopted here because, following de Finetti, “coherence” is used more broadly to refer to preferences that are consistent in the sense of not leading to arbitrage.

One more technical detail is needed, namely a strengthening of axiom 4. Define an allocation $\boldsymbol{w}$ to be an *H-bet* if it is constant within each element of *H*.

Axiom 4*: (Conditional overlap among strategies) There exist *H*-bets { $\boldsymbol{w}^{(a)}$, $a \in A$ } in the interior of *W* such that $(a,\ \boldsymbol{w}^{(a)}) \sim_E (b,\ \boldsymbol{w}^{(b)})$ for every event $E \subseteq I$ and all pure strategies *a* and *b*.

Axiom 4* requires that differences among pure strategies can be fully compensated not only unconditionally but also conditionally on elements of *I*, via payments that depend (only) on events in *H*, which is reasonable if the decision maker has intrinsic interests or prior stakes in the hypotheses but not in the informational events. Significantly, though, there are no constant *acts* here, nor any comparisons between acts conditioned on two different events.

Theorem 2: Axioms 1-2-3-4*-5-6 are satisfied by $\succcurlyeq$ and all $\succcurlyeq_E$ if and only if there exist strictly increasing continuous evaluation functions for money $\{v_{as}\}$, such that $\succcurlyeq$ is represented by a utility function *U* having the additive form (1) and for any event $E \subseteq I$ the conditional preference relation $\succcurlyeq_E$ is represented by

$$U_{\tilde{\boldsymbol{a}}}(\boldsymbol{w} \mid E) = \sum_{s \in E} v_{\tilde{a}(s)s}(w_s), \qquad (2)$$

with $U_a(\boldsymbol{w}^{(a)} | E) = U_b(\boldsymbol{w}^{(b)} | E)$ for all pure strategies *a* and *b*. *U* is unique up to positive affine transformations, and $\{v_{as}\}$ are unique up to joint transformations of the form $\alpha v_{as} + \beta_s + \gamma_{as}$

where α is an arbitrary positive scale factor and $\{\beta_s\}$ and $\{\gamma_{as}\}$ are arbitrary constants with $\{\gamma_{as}\}$ summing to zero within each element of I for every a.

Thus, conditional utility is obtained from unconditional utility merely by restricting the summation of evaluation functions to the appropriate subset of states, which of course is consistent with the usual Bayesian updating rule when preferences have a subjective expected utility representation. Conversely, unconditional utility can be expressed as the sum of conditional utilities over a set of mutually exclusive and exhaustive events. In particular, the unconditional utility of $(\tilde{\boldsymbol{a}}, \boldsymbol{w})$ is:

$$U_{\tilde{\boldsymbol{a}}}(\boldsymbol{w}) = \sum_{k=1}^{K} U_{\tilde{\boldsymbol{a}}}(\boldsymbol{w} \mid I_k)$$

where $U_{\tilde{\boldsymbol{a}}}(\boldsymbol{w} \mid I_k) = U_{\boldsymbol{a}}(\boldsymbol{w} \mid I_k)$ for the pure strategy a that agrees with $\tilde{\boldsymbol{a}}$ on I_k.

4. Elicitation of Preference Parameters Through Belief and Value Gambles

Given a theoretical representation of the decision maker’s preferences on a large set of both real and hypothetical acts, such as a subjective expected utility model or the more general additive model of this paper, there remains the question of how its parameters might be elicited and credibly communicated in a concrete economic setting where only a small subset of acts are really available. If other actors are present in the same scene—say, decision analysts or principals or opposing players—the question of who-knows-what-about-whom-and-how-do-they-know-it becomes important.

In the setting of Bayesian decision analysis, it is often assumed that the decision maker is not fully aware of all her preferences at the outset: she may have to *construct* them by pondering the situation from multiple perspectives, decomposing it into smaller parts, comparing simple alternatives and using normative axioms to reconcile inconsistencies and extrapolate to more complex alternatives. In Savage's own words, this may involve comparisons among counterfactual acts such as "being hung without damage to one's health or reputation." In practice, these comparisons are usually designed to separately measure probabilities and utilities by imagining acts that yield only one or two distinct consequences and acts that yield objective lotteries over consequences. There are two potential problems with this approach. One is that the subjective probabilities and utilities determined by such measurements may not reflect the decision maker's true beliefs and values if her utilities for consequences are in fact state-dependent. (The possibility of state-dependent utility scale factors can never be ruled out, even if the usual axiom of state-independence—Savage's P3—is satisfied.) The other problem is that preference assertions that refer to counterfactual acts (e.g., constant acts or objective lotteries over subjective consequences) may not be credible to observers if they cannot be validated by material choices that are really available: they are merely unsubstantiated verbal expressions.

In the setting of a noncooperative game, it is typically assumed that the players start by knowing their own—and everyone else's—utilities for outcomes, but not their probability distributions concerning each other's strategies, and the goal is to endogeneously determine those probabilities according to an appropriate concept of interactive rationality. Here, the problem is not only how to separate the players' utilities from their probabilities via simple and credible measurements, but how to do so in a way that renders the results *common knowledge* in an interactive setting: how does one player know that what she has revealed about her preferences has been heard or believed by her opponent, and vice versa ad infinitum? And what is the appropriate standard of interactive rationality from a Bayesian viewpoint? Do Savage's axioms even apply when the "states" perceived by one player are "strategies" under the control of an intelligent opponent? There is still no consensus on these issues. (Aumann and Drèze 2004 offer a new approach that generalizes the Anscombe-Aumann model to strategic risk.)

The general additive model circumvents the difficulties just mentioned insofar as (i) it does not emphasize the separation of probabilities from state-dependent utilities, (ii) it does not involve acts with unnatural mappings of consequences to states, (iii) it treats "moves" of the decision maker and the opponent symmetrically, providing an immediate path of generalization from decisions to games, and (iv) it gives an explicit role to money, which provides a language for credible public communication (namely the language of contingent claim contracts) as well as an operational standard of interactive rationality (namely, that no arbitrage profits should be realized by an observer). The additive model implies that a decision maker who wishes her preferences to be known can reveal them through the public acceptance of small gambles, in the manner of de Finetti's (1937, 1974) operational definition of subjective probability. While de Finetti's gambling method is usually presented as an elegant and simple way to derive the laws of probability from a no-arbitrage argument (under a seemingly-restrictive assumption of linear utility for money), it can also be adapted to the measurement of values as well as beliefs, and it has an another, under-appreciated virtue, namely that a public offer by the decision maker to accept a gamble not only reveals information about her beliefs and values, but it does so in a way that renders the information common knowledge, in the sense of the "specular" common knowledge of a public market (Nau 1995b).

Henceforth, assume that the concrete choice problem facing the decision maker includes only a finite number of feasible acts (strategy-allocation pairs), together with the possibility of small side bets, and without loss of generality let each feasible act be identified with a pure strategy with zero additional allocation of wealth. Thus, "strategy *a*" is to be henceforth interpreted as “act (a, $\mathbf{0}$)”. Whereas traditional models of decisions and games strive for a complete separation of the decision maker’s beliefs from her values, the former represented by probabilities and the latter by utilities, the achievement of that goal is generally impossible under the assumptions of the general additive model. Nevertheless, a partial separation can be achieved through the articulation of two distinct kinds of gambles:

(i) [given my present information] “if strategy *a* were chosen, then I would accept the following gambles on states…” and

(ii) “if I were to choose strategy *a* when *b* is also available [regardless of my information], then I would also accept the following gambles on states…”

From the viewpoint of an observer, these are conditional bets on a larger state space, namely $A \times S$, in which the conditioning events happen to be the elements of *A*, the decision maker’s possible strategy choices. However, very importantly, the decision maker is not asked to bet directly *on* her own strategy choices. Type (i) gambles will be called *belief gambles* because they reveal information about present beliefs, in the manner of de Finetti, although those beliefs may be distorted by state- and strategy-dependent marginal utilities for money. Type (ii) gambles will be called *value gambles* because they reveal information about differences in profiles of state-dependent values between pairs of acts, although those values too may be distorted by marginal utilities for money. In fact, it will be seen that the distortions of belief and value gambles by unobservable marginal utilities for money are precisely reciprocal to each other and hence not catastrophic for decision analysis.[2]

[2] A value gamble is a gamble that would be acceptable even if beliefs were updated by new information, and thus it is independent of beliefs. A belief gamble may reflect either the decision maker’s present unconditional beliefs (assuming that no more information will arrive before a strategy is chosen) or else the conditional beliefs that she

First, consider the construction of belief gambles. In the event where strategy a is chosen (for whatever reason), what small gambles should the decision maker accept if she doesn't mind publicly revealing her preferences (or if she believes they are already public knowledge)? Evidently she should accept any small gamble $\boldsymbol{z}$ such that $U_a(\boldsymbol{z}) \geq U_a(\mathbf{0})$. Assuming differentiability[3] of the evaluation functions for money, the first order condition is that $\boldsymbol{z}$ should have non-negative expectation with respect to a probability distribution $\boldsymbol{\pi}_a(\mathbf{0})$ determined by evaluating

$$\pi_{as}(\boldsymbol{w}) = \frac{v'_{as}(w_s)}{\sum_{r \in S} v'_{ar}(w_r)} \tag{3}$$

at $\boldsymbol{w} = \mathbf{0}$, where v'_{as} denotes the first derivative of v_{as}. The distribution $\boldsymbol{\pi}_a(\boldsymbol{w})$ is the decision maker's *risk neutral distribution at allocation* $\boldsymbol{w}$ *under strategy a*, since under those conditions she accepts small gambles as though she were risk neutral with that probability distribution. The dependence of $\boldsymbol{\pi}_a(\boldsymbol{w})$ on $\boldsymbol{w}$ encodes the decision maker's attitude toward risk when strategy a is chosen, as described in Nau (2003, 2005). The fact that the utility function is additively separable means that the decision maker is *uncertainty neutral* although she may be risk averse. For a given strategy a, the distribution $\boldsymbol{\pi}_a(\mathbf{0})$ can be fully revealed by the acceptance of a finite number of belief gambles—in fact, $2(N-1)$ one-sided gambles, where $N=|S|$ is the number of states. This is merely de Finetti's operational method of measuring probabilities, tempered by the recognition that the probabilities revealed need not be the decision maker's true subjective

expects to hold in the event a particular strategy is chosen (as if she is awaiting a signal or concealing private information). In a game situation, it would be natural for value gambles (the "rules of the game") to be articulated prior to belief gambles, to allow for endogenous belief formation. In a decision analysis situation, belief gambles might be assessed first, before values are fully structured..

[3] Henceforth it will be assumed that preferences are sufficiently smooth so that the evaluation functions $\{v_{as}\}$are continuously differentiable. Alternatively, it could be assumed that preferences among allocations are convex for every strategy, which would yield differentiability almost everywhere, as well as risk aversion.

probabilities. Under the general additive model, the decision maker's probabilities are entirely undefined, although her marginal betting rates for money are well-defined given any act.

Next, consider the construction of value gambles. In the event where strategy a is chosen when b is available, what small gambles should the decision maker accept (if she doesn't mind revealing her preferences, etc.)? The choice of a over b evidently means that she ought to accept a gamble $\boldsymbol{\mu}_{ab}$ whose state-by-state increments of cardinal utility are proportional to the state-by-state *differences* in utility between a and b, evaluated from the perspective of the local marginal utilities for money that apply under the chosen strategy a. This implies that $\boldsymbol{\mu}_{ab}$ is the vector whose s^{th} element is:

$$\mu_{abs} \propto \frac{v_{as}(0) - v_{bs}(0)}{v'_{as}(0)} \qquad (4)$$

The direct contemplation of such a gamble is admittedly more subtle than a simple bet on states given a, but the key idea is that the comparison of a-plus-a-small-side-bet-proportional-to-$\boldsymbol{\mu}_{ab}$ against a alone should "feel" just like a microcosm of the comparison of a against b in the sense that it yields a proportional profile of changes in utility across states. A more decomposed method of constructing value gambles is described by Nau (1995a). Note that, if the decision maker is at heart a state-dependent expected utility maximizer with strategy-independent beliefs, so that $v_{as}(w) = p_s u_{as}(w)$, then her acceptable belief gambles depend on both her probabilities and her utilities, because her risk neutral probabilities satisfy $\pi_{as}(\boldsymbol{w}) \propto p_s u'_{as}(w)$ for each a, while her acceptable value gambles depend on her utilities alone, because p_s drops out of (4). However, it is behaviorally impossible to distinguish this case from the case in which $v_{as}(w) = p_{as}u_{as}(w)$, where there is some dependence of probabilities on strategies that would play a role in determining the acceptable value gambles.

The formal rules of the gambling game are as follows. For each available strategy, let the decision maker articulate a complete set of acceptable belief gambles and complete set of acceptable value gambles. A complete set of belief gambles for a given strategy is a set of gambles sufficient to uniquely determine a risk neutral probability distribution. A complete set of value gambles for a given strategy includes gambles comparing that strategy against all other available strategies. If M is the number of available strategies, then in total $2M(N-1)$ belief gambles and $M(M-1)$ value gambles are required. Each belief and value gamble is conditioned on a particular strategy and yields non-zero payoffs only if that strategy is chosen. In general, for the decision maker to say that "gamble $\mathbf{z}$ is acceptable in the event that strategy a is chosen," where $\mathbf{z}$ denotes a vector of (positive and negative) monetary payoffs over states, means that for any sufficiently small[4] non-negative coefficient λ chosen by a betting opponent, the decision maker is willing to receive the payoff $\lambda\mathbf{z}$ if a is chosen, and zero otherwise, thereby experiencing the act $(a, \lambda\mathbf{z})$ instead of $(a, \mathbf{0})$. Furthermore, under the first-order utility approximation, the *sum* of acceptable gambles is also acceptable (utility-non-decreasing). Hence the betting opponent may assign a separate non-negative coefficient to each gamble that is accepted, and the net transaction will be a corresponding linear combination of their payoffs. The entire sequence of events then unfolds like this: (i) the decision maker and a betting opponent perceive a finite set A of available strategies and a finite set S of possible states, (ii) the decision maker reveals information about her preferences in the form of belief and value gambles [in either order], (iii) the opponent chooses small non-negative multipliers for the gambles, (iv) the decision maker

[4] If the evaluation functions are nonlinear (e.g., risk averse), a slight technical modification is needed, namely a concept of "ε-acceptability" of a gamble $\mathbf{z}$ in which for every $\varepsilon > 0$ there exists $\delta > 0$ such that for any non-negative $\lambda < \delta$ chosen by a betting opponent, the decision maker will agree to receive the payoff $\lambda(\mathbf{z}+\varepsilon)$ if a is chosen. The definition of an arbitrage opportunity is generalized accordingly. See Nau (1995a) for details.

chooses one of the available strategies, (v) the state of the world is revealed and (vi) the gambles are settled.

By multiplying (3) and (4) and summing across states, it can seen that $U_a(\mathbf{0}) \geq U_b(\mathbf{0})$ if and only if $\boldsymbol{\pi}_a(\mathbf{0}) \cdot \boldsymbol{\mu}_{ab} \geq 0$, so that the revealed belief and value gambles suffice to determine the decision maker's ordering of the available strategies.[5] However, there is another, deeper sense in which the revealed gambles determine the decision maker's "Bayesian rational" choice. Suppose the betting opponent is an arbitrageur who is only interested in riskless profit opportunities, i.e., Dutch books. By linear duality, it can shown that the opponent is able to achieve an *ex post arbitrage profit* if and only if the decision maker fails to choose the utility-maximizing act (Nau 1995a). That is, the opponent can find a combination of gambles such that he will not lose money in any event but he will win some money if the decision maker fails to choose the utility-maximizing act. From the opponent's perspective, it is as if the decision maker has assigned probability zero to the event that she will choose a strategy that fails to maximize her utility, in the sense that she is willing to throw money away if such a strategy is chosen. Thus, belief and value gambles provide a basis for "decision analysis by arbitrage," and because no-arbitrage also characterizes group rationality, this solution concept extends without modification to noncooperative games, where it is the dual definition of *correlated equilibrium*.[6]

---

[5] Let $\boldsymbol{v}_a$ denote the evaluation vector whose s^{th} element is $v_{as}(\mathbf{0})$, whose elements sum to $U_a(\mathbf{0})$. Then $\boldsymbol{\pi}_a(\mathbf{0}) \times \boldsymbol{\mu}_{ab} \propto \boldsymbol{v}_a - \boldsymbol{v}_b$, hence $\boldsymbol{\pi}_a(\mathbf{0}) \cdot \boldsymbol{\mu}_{ab} \propto U_a(\mathbf{0}) - U_b(\mathbf{0})$.

[6] In a noncooperative game, there are two or more decision makers, and one decision maker's strategies correspond to another's states of the world, and vice versa. In this setting, if the players publicly reveal information about their values for outcomes through the medium of value gambles (leaving beliefs to be endogenously determined), then the observer is unable to reap an arbitrage profit if and only if the outcome of the game is one that has positive probability in an appropriate form of correlated equilibrium. If the players have constant marginal utility for money (quasilinear utility), the supporting equilibrium is an objective correlated equilibrium (Nau & McCardle 1990); otherwise it is a subjective correlated equilibrium with common prior risk-neutral probabilities (Nau 1995c). In a game of incomplete information (where belief gambles also play a role in articulating a common prior over types), the correlated equilibrium is a generalized Bayesian or communication equilibrium (Nau 1992). Hence the general additive model (1) provides sufficient structure on preferences for the modeling of noncooperative games—without

5. BAYESIAN UPDATING

The Bayesian properties of the general additive model will now be explored. As in §3, let the partitioned state space be written as $S = H \times I = \{H_1, \ldots H_J\} \times \{I_1, \ldots, I_K\}$ where $\{H_j\}$ are hypotheses of interest to the decision maker and $\{I_k\}$ are informational results of an experiment on which acts and preferences may be conditioned, and henceforth let $\succcurlyeq_k$ denote the decision maker's conditional preference relation given event I_k. States of the world consist of products H_jI_k and are indexed by jk, in terms of which the general additive model can be re-expressed as:

$$U_a(\boldsymbol{w}) = \sum_{j=1}^{J}\sum_{k=1}^{K} v_{ajk}(w_{jk})\,.$$

Conditional on ("posterior to") observing the result I_k, the decision maker's preferences are represented by:

$$U_a(\boldsymbol{w} \mid I_k) = \sum_{j=1}^{J} v_{ajk}(w_{jk})\,, \tag{5}$$

in which a particular k rather than a summation over all k appears. Now assume that, as in most scientific applications, the decision maker has no intrinsic interest in the experimental result given the hypothesis, so that her utility of money does not intrinsically depend on k, and moreover she believes that the experimental result does not intrinsically depend on her strategy given the hypothesis. (This does not rule out the possibility that her strategy could somehow affect the probability of the hypothesis, which is behaviorally indistinguishable from utility with strategy-hypothesis interactions.) For any allocation $\boldsymbol{y}$, let $w_j\,\boldsymbol{y}_{-j}$ denote the allocation that yields

necessarily requiring the players to know each others' true utilities or agree on true probabilities—and correlated equilibrium emerges as the natural game-theoretic solution concept under the no-arbitrage standard of Bayesian rationality. Aumann (1987) has previously argued that correlated equilibrium is the natural Bayesian solution concept, although he leaves open the question of how the players' utilities become common knowledge.

payoff w in hypothesis j (for all k) and agrees with $\boldsymbol{y}$ elsewhere. The preceding assumptions are formalized in:

Axiom 7: (No intrinsic interest or influence over the experiment, given the hypothesis) For every hypothesis j, pair of strategies a and b, allocations $\boldsymbol{y}$ and $\boldsymbol{z}$, and constants w, x, w^* and x^*: $(a, w_j\, \boldsymbol{y}_{-j}) \preccurlyeq_k (b, x_j\, \boldsymbol{z}_{-j})$ and $(a, w^*_j\, \boldsymbol{y}_{-j}) \succcurlyeq_k (b, x^*_j\, \boldsymbol{z}_{-j})$ and $(a, w_j\, \boldsymbol{y}_{-j}) \succcurlyeq_m (b, x_j\, \boldsymbol{z}_{-j}) \Rightarrow (a, w^*_j\, \boldsymbol{y}_{-j}) \succcurlyeq_m (b, x^*_j\, \boldsymbol{z}_{-j})$.

In words, *if* changing the payoff from w to w^* in hypothesis j under strategy a is no worse than changing the payoff from x to x^* in hypothesis j under strategy b, conditional on result k, *then* changing w to w^* in hypothesis j under strategy a cannot be strictly worse than changing x to x^* in hypothesis j under strategy b conditional on some *other* result m, ceteris paribus. When added to the other six axioms, Axiom 7 implies that, for a given strategy and hypothesis, the evaluation functions for money have the same shape under all outcomes of the experiment, but they may have different scale factors, and the dependence of the scale factors on the experimental results is the same under every strategy. With suitable normalization, the scale factors can be viewed as conditional probabilities of experimental results given hypotheses. This result is formalized as:

Theorem 3: Axiom 7 is satisfied in addition to 1-2-3-4*-5-6 if and only if the evaluation functions in the conditional utility function (2) satisfy

$$v_{ajk}(w) = p_{k|j}\, u_{aj}(w) + \beta_{jk}$$

where $\{u_{aj}\}$ are strictly increasing continuous functions, $p_{k|j} > 0$, $\sum_{k=1}^{K} p_{k|j} = 1$ for each j, $\{\beta_{jk}\}$ are constants and $u_{aj}(w_j^{(a)}) = u_{bj}(w_j^{(b)})$ for all a, b, j.

Here $u_{aj}(w)$ may be interpreted as the product of the decision maker's possibly-strategy-dependent prior probability for hypothesis j and possibly-strategy-and-hypothesis-dependent utility for money, and $p_{k|j}$ can be interpreted as a strategy-*independent* conditional probability for result k given hypothesis j, i.e., the decision maker's likelihood function for the experiment.

Now consider the application of the decision analysis method of the preceding section. As a purely formal exercise, the risk neutral distribution under strategy a can always be uniquely factored as the product of a marginal (prior) distribution over hypotheses, $\boldsymbol{\theta}_a$, and a conditional distribution (likelihood) over experimental results given the hypothesis, $\phi_{ak|j}$:

$$\pi_{ajk}(\mathbf{0}) = \theta_{aj}\phi_{ak|j}.$$

Under the representation of Theorem 3, it is straightforward to show that the risk-neutral likelihood is independent of the strategy and is identical to the "true" likelihood, i.e., $\phi_{ak|j} = p_{k|j}$ for all available strategies. Furthermore, it is straightforward to show that for every value gamble, the payoff depends only on the hypothesis, not on the experimental result. Thus,

$$\mu_{abjk} \equiv \hat{\mu}_{abj}$$

where $\hat{\boldsymbol{\mu}}_{ab}$ is a J-vector representing a reduced form value gamble over hypotheses alone. These two pieces of evidence (risk neutral likelihood independent of the strategy, value gamble payoff independent of the experimental result) reveal that it is "as if" the decision maker satisfies Axiom 7, and they allow an unbiased assessment of the likelihood function that she ascribes to the experiment. However, the decision maker's prior distribution is not revealed, being entangled with her marginal utilities that may depend on the strategy as well as on the hypothesis, so that in general $\boldsymbol{\theta}_a \neq \boldsymbol{\theta}_b$ for $a \neq b$.

Under these conditions, by applying the decision analysis procedure of the preceding section, the decision maker's prior (unconditional) preferences among strategies are completely described by her (strategy-dependent) risk neutral prior distributions $\{\boldsymbol{\theta}_a\}$ and her reduced-form value gambles $\{\hat{\boldsymbol{\mu}}_{ab}\}$. In particular, strategy a is preferred to b iff $\boldsymbol{\theta}_a \cdot \hat{\boldsymbol{\mu}}_{ab} \geq 0$. Now consider the updating of preferences after observing result I_k. The decision maker's acceptable value gambles under the conditioned utility function (2) are identical to the reduced-form value gambles from the prior stage. (The redundant elements in the original value gambles are merely compressed out.) Her strategy-dependent posterior risk neutral distributions are obtained by applying Bayes' rule to the prior distributions using the common likelihood: the posterior distribution under strategy a given I_k is the vector $\bar{\boldsymbol{\theta}}_{ak}$ whose j^{th} element is $\bar{\theta}_{akj} \propto \theta_{aj} p_{k|j}$. Finally, the decision maker prefers a to b posterior to observing I_k iff $\bar{\boldsymbol{\theta}}_{ak} \cdot \hat{\boldsymbol{\mu}}_{ab} \geq 0$.

Thus, under conditions where the decision maker has no intrinsic interest or influence over the experiment given the hypothesis, Bayesian updating applies to the decision analysis procedure of the preceding section, with the (strategy-dependent) *risk neutral* prior distributions being updated through multiplication by an observable (and strategy-independent) likelihood function. The fact that the likelihood function is observable makes it reasonable to use "objective" likelihoods where data or theory are applicable and also makes it plausible for different individuals to agree on the likelihood. The decision maker's prior distribution is generally private and unobservable due to entanglement with state-and-strategy-dependent utility and personal stakes in the hypotheses, but this does not matter: the distorting effects of state-dependent utilities are merely lumped together with other private sources of subjectivity. The general additive utility model therefore supports statistical procedures that are based on the likelihood principle and provides a foundation on which Bayesian inference and decision analysis can be carried out without introducing Savage-type consequences, without referring to counterfactual acts in the process of belief and value elicitation, and without uniquely separating prior probabilities from utilities. All that is required is an additive representation of preferences over the observable "moves" of the decision maker and the adversary (nature or whomever), with money serving as a medium of exchange and credible communication.

6. DISCUSSION

Recently Machina (2003, 2004) and Karni (2005) have criticized the traditional choice-theoretic concept of states of the world. Machina points out that, in Savage's SEU theory, states of the world are "endogeneous [to] and constructed [by]" the decision maker. Motivated in part by Machina's observations, Karni presents an axiomatic model of "SEU without states of the world" in which the primitives are "actions" that may be undertaken and "effects" that may be experienced by the decision maker. The actions and effects are assumed to be observable in the sense that monetary bets can be placed on them. (In this respect Karni's model resembles the model of this paper, except that here bets are placed on combinations of actions and states rather than actions and effects.) Karni's axioms imply that preferences are uniquely represented by action-independent utilities for effects combined with action-dependent probabilities for effects. However, it does not necessarily follow that the probabilities in that representation are the decision maker's "true" probabilities: there are other equivalent representations in which different probabilities are combined with action-dependent utilities.

The point of view expressed in this paper is that the problem with Savage's framework lies not in the concept of states of the world but rather in the concept of consequences as states-of-the-person that can be assigned to arbitrary states-of-the-world without affecting their utilities, all in pursuit of an elusive definition of true subjective probabilities. The states of the *person*, rather than states of the world, are what are "endogenous and constructed." By abandoning the personalistic concept of consequences and the goal of extracting true subjective probabilities, and by relying instead on money as a yardstick for measurements, it becomes possible to redefine the states of the world so that they refer (only) to events on which bets can be placed and whose realization can be agreed upon by different observers. The set of states of the world can then be as coarse as needed to meet scientific standards of exogeneity and observability.

REFERENCES

[1] ANSCOMBE, F. AND R. AUMANN (1963) A Definition of Subjective Probability. *Annals of Mathematical Statistics* **34**, 199—205.

[2] AUMANN, R. (1971) Letter to L.J. Savage. Reprinted in Drèze, J. (1987) *Essays on Economic Decisions under Uncertainty*. Cambridge Univ. Press, London.

[3] AUMANN, R. (1987) Correlated Equilibrium as an Expression of Bayesian Rationality. *Econometrica*. **55***,* 1—18

[4] AUMANN, R. AND J. DRÈZE (2004) Assessing Strategic Risk. Working paper, Center for Rationality, Hebrew University of Jerusalem

[5] DE FINETTI, B. (1937) La Prévision: Ses Lois Logiques, Ses Sources Subjectives. *Ann. Inst. Henri Poincaré* **7,** 1—68. Translation reprinted in H.E. Kyburg and H.E. Smokler, Eds. (1980), *Studies in Subjective Probability*, 2nd ed., Robert Krieger, New York, 53—118.

[6] DE FINETTI, B. (1974), *Theory of Probability*, Vol. 1. New York: John Wiley and Sons

[7] KARNI, E. (2004) Foundations of Bayesian Theory. Working paper, Johns Hopkins University

[8] KARNI, E. (2005) Subjective Expected Utility Without States of the World. Working paper, Johns Hopkins University

[9] KARNI, E., AND D. SCHMEIDLER (1980) An Expected Utility Theory for State-Dependent Preferences. Working paper 48-80. Foerder Institute for Economic Research, Tel-Aviv University

[10] KARNI, E. AND P. MONGIN (2000) On the Determination of Subjective Probability by Choices. *Management Science* **46**, 233—248.

[11] KARNI, E., D. SCHMEIDLER AND K. VIND (1983) On State Dependent Preferences and Subjective Probabilities. *Econometrica* **51**, 1021—1031.

[12] MACHINA, M. (2003) States of the World and the State of Decision Theory. In Donald Meyer (ed.) *The Economics of Risk*, W. E. Upjohn Institute

[13] MACHINA, M. (2004) Almost-Objective Uncertainty. *Economic Theory* **24**, 1—54.

[14] NAU, R. (1992) Joint Coherence in Games of Incomplete Information. *Management Science* **38**:3, 374—387

[15] NAU, R. (1995a) Coherent Decision Analysis with Inseparable Probabilities and Utilities. *J. Risk and Uncertainty* **10**, 71—91.

[16] NAU, R. (1995b) The Incoherence of Agreeing to Disagree. *Theory and Decision* **39**, 219—239

[17] NAU, R. (1995c) Arbitrage-Free Correlated Equilibria. Working Paper, Fuqua School of Business

[18] NAU, R. (2001) De Finetti Was Right: Probability Does Not Exist. *Theory and Decision* **59**, 89—124.

[19] NAU, R. (2003) A Generalization Of Pratt-Arrow Measure To Non-Expected-Utility Preferences And Inseparable Probability And Utility. *Management Science* **49**:8, 1089—1104

[20] NAU, R. (2005) Uncertainty Aversion with Second-Order Utilities and Probabilities. *Management Science*, forthcoming

[21] NAU, R. AND K. MCCARDLE (1990) Coherent Behavior in Noncooperative Games. *Journal of Economic Theory* **50**:2, 424—444

[22] NAU, R. AND K. MCCARDLE (1991) Arbitrage, Rationality, and Equilibrium. *Theory and Decision* **31,** 199—240

[23] SAVAGE, L. (1954) *The Foundations of Statistics*. Wiley, New York. (Second edition 1972, Dover, New York)

[24] SCHERVISH, M., T. SEIDENFELD, AND J. KADANE (1990) State-Dependent Utilities. *J. American Statistical Association* **85**, 840–847.

[25] SHAFER, G. (1986) Savage Revisited. *Statistical Science* **1**, 463—501.

[26] SKIADAS, C. (1997a) Conditioning and Aggregation of Preferences. *Econometrica* **65**: 2, 347—367

[27] SKIADAS, C. (1997b) Subjective Probability under Additive Aggregation of Preferences. *J. Econ. Theory* **76**, 242—271

[28] WAKKER, P. (1989) Additive Representations of Preferences: A New Foundation of Decision Analysis. Kluwer Academic Publishers, Dordrecht.

APPENDIX: SKETCHES OF PROOFS

Theorem 1: By axioms 1-3, for a fixed strategy a, $\succcurlyeq$ is represented by an additively separable utility function:

$$U_a(\boldsymbol{w}) = \sum_{s\in S} v_{as}(w_s)$$

that is unique up to positive affine transformations. This follows from standard arguments (e.g., Theorem III.6.6 in Wakker 1989) since axiom 3 reduces to ordinary generalized triple cancellation when b=a. Next, for strategy a, form a standard sequence of w's that are equally spaced in terms of cardinal a-utility, then form a sequence of x's that are indifferent to that sequence of w's when received in conjunction with any other strategy b. Then use axiom 3 to show that these x's are equally spaced in terms of cardinal b-utility, which fixes the units of a-utility and b-utility relative to each other. Finally, use axiom 4 to fix the origins of the utility scales relative to each other so that $U_a(\boldsymbol{w}^{(a)}) = U_b(\boldsymbol{w}^{(b)})$.

Details: By axiom 4 there exist interior $\boldsymbol{w}$ and $\boldsymbol{x}$ such that $(a, \boldsymbol{w}) \sim (b, \boldsymbol{x})$. By axiom 2, there exist $w^*_s > w_s$ and $x^*_s > x_s$ such that $(a, w^*_s\boldsymbol{w}_{-s}) \sim (b, x^*_s\, \boldsymbol{x}_{-s})$, as well as $\boldsymbol{w}^{**}_{-s}$ and $\boldsymbol{x}^{**}_{-s}$ such that $(a, w_s\boldsymbol{w}^{**}_{-s}) \sim (a, w^*_s\, \boldsymbol{w}_{-s}) \sim (b, x^*_s\, \boldsymbol{x}_{-s}) \sim (b, x_s\, \boldsymbol{x}^{**}_{-s})$. Now consider $(a, w^*_s\, \boldsymbol{w}^{**}_{-s})$ and $(b, x^*_s\boldsymbol{x}^{**}_{-s})$. The increment of a-utility in going from $(a, \boldsymbol{w})$ to $(a, w^*_s\, \boldsymbol{w}_{-s})$ must be the same as in going from $(a, w^*_s\, \boldsymbol{w}_{-s})$ to $(a, w^*_s\, \boldsymbol{w}^{**}_{-s})$ because this is one step in a standard sequence. Similarly, the increment of b-utility in going from $(b, \boldsymbol{x})$ to $(b, x^*_s\, \boldsymbol{x}_{-s})$ must be the same as in going from $(b, x^*_s\, \boldsymbol{x}_{-s})$ to $(b, x^*_s\, \boldsymbol{x}^{**}_{-s})$. But also, by axiom 3, $(a, w^*_s\, \boldsymbol{w}^{**}_{-s}) \sim (b, x^*_s\, \boldsymbol{x}^{**}_{-s})$. Therefore, increments of a-utility are proportional to increments of b-utility. Without loss of generality, the origin and scale of $U_b(.)$ may be chosen so that $U_a(\boldsymbol{w}) = U_b(\boldsymbol{x})$ and so that the constant of proportionality between a-utiles and b-utiles is equal to 1. Next find a third strategy

c and repeat the process to fix the origin and scale of $U_c(.)$, and so on until all the strategies have been covered. ■

Theorem 2: Under axioms 1-3 and 4*, the representation of Theorem 1 holds, and without loss of generality, the constants in the evaluation functions can be chosen so that $\{v_{as}(w_s^{(a)})\}$ sum to zero within each element of I for every a. By axiom 5, the conditional relation $(a, \boldsymbol{w}) \succcurlyeq_E (b, \boldsymbol{x})$ depends only on $\boldsymbol{w}_E$ and $\boldsymbol{x}_E$. Therefore, let $\boldsymbol{w}_{-E} = \boldsymbol{w}_{-E}^{(a)}$ and $\boldsymbol{x}_{-E} = \boldsymbol{w}_{-E}^{(b)}$. By axiom 6, letting $F = S \backslash E$, it follows that $(a, \boldsymbol{w}) \succcurlyeq_E (b, \boldsymbol{x})$ iff $(a, \boldsymbol{w}_E \boldsymbol{w}_{-E}^{(a)}) \succcurlyeq (b, \boldsymbol{x}_E\ \boldsymbol{w}_{-E}^{(b)})$, which (under the representation of Theorem 1) is true if and only if

$$\sum_{s \in E} v_{as}(w_s) \geq \sum_{s \in E} v_{bs}(x_s)$$

which establishes (2) as the representation of $\succcurlyeq_E$.■

Theorem 3: First, starting from the conditional utility representation (2) implied by axioms 1-6, apply axiom 7 with $a = b$ and $\boldsymbol{y} = \boldsymbol{z}$ to show that $v_{ajk}(w^*) - v_{ajk}(w) \geq v_{ajk}(x^*) - v_{ajk}(x) \Rightarrow v_{ajm}(w^*) - v_{ajm}(w) \geq v_{ajm}(x^*) - v_{ajm}(x)$, which implies that a sequence of payoffs that are equally spaced in terms of ajk-utility are also equally spaced in terms of ajm-utility, hence $v_{ajk}(w) = c_{ajk}\, u_{aj}(w) + \beta_{ajk}$ for some positive constants $\{c_{ajk}\}$, constants $\{\beta_{ajk}\}$, and functions $\{u_{aj}\}$. Second, apply axiom 7 with $a \neq b$, $\boldsymbol{y} = \boldsymbol{z}$, to show that $\beta_{ajk} \equiv \beta_{jk}$ independent of the strategy a. Third, apply axiom 7 with $a \neq b$, $\boldsymbol{y} = \boldsymbol{w}^{(a)}$, and $\boldsymbol{z} = \boldsymbol{w}^{(b)}$, $w = w_{jk}^{(a)} = w_{jm}^{(a)}$ and $x = w_{jk}^{(b)} = w_{jm}^{(b)}$ to show that $v_{ajk}(w^*) - v_{ajk}(w) \geq v_{bjk}(x^*) - v_{bjk}(x) \Rightarrow v_{ajm}(w^*) - v_{ajm}(w) \geq v_{bjm}(x^*) - v_{bjm}(x)$, which implies that $c_{ajk}/c_{ajm} = c_{bjk}/c_{bjm}$, whence by suitable normalization of the functions u_{aj} and u_{bj}, it is possible to force $c_{ajk} = c_{bjk} \equiv p_{k|j}$, with $\sum_{k=1}^{K} p_{k|j} = 1$, $u_{aj}(w_j^{(a)}) = u_{bj}(w_j^{(b)})$ for all a, b, j, and $\{\gamma_{ajk}\}$ summing to zero over k. ■