

Subjective Probability*

Nabil I. Al-Najjar[†] and Luciano De Castro[‡]

Northwestern University

March 2010

Abstract

We provide an overview of the idea of subjective probability and its foundational role in decision making and modern management sciences. We highlight the role of Savage's theory as an organizing methodology to guide and constrain our modeling of choice under uncertainty, rather than a substantive statement subject to refutations by experimental or psychological evidence.

*Prepared for *The Wiley Encyclopedia of Operations Research and Management Science*.

[†] Department of Managerial Economics and Decision Sciences, Kellogg School of Management, Northwestern University, Evanston IL 60208.

[‡] Department of Managerial Economics and Decision Sciences, Kellogg School of Management, Northwestern University, Evanston IL 60208. At the University of Illinois, Urbana-Champaign until June 2010.

Contents

1	Introduction	1
2	Expected Utility Theory	3
2.1	Von Neumann-Morgenstern Representation	3
2.2	Savage's Framework	4
2.3	Savage's Axioms	5
2.4	Savage's Subjective Expected Utility Representation	8
3	Interpretation and Implications	10
3.1	Normative Theory and the Definition of 'Rationality'	10
3.2	Feasibility Constraints and Complexity	13
3.3	Information, Dynamic Choice, and Dynamic Consistency	14
3.4	Risk vs. Uncertainty	15
3.5	Exchangeability, Objectivity, and Frequencies	18
4	Concluding Remarks	20

1 Introduction

What is probability? Observers of scientific progress in the last several decades will likely find this question puzzling. Modern probability theory is, by all objective measures, a runaway success in shaping modern science. In management sciences, entire fields, such as finance, economics, and operations research are in part founded on probabilistic concepts and tools. Yet it is hard to think of other concepts as important as probability whose very meaning remains unclear and, often, controversial.

The formative years of the modern theory of probability, roughly from the 1920's through the 1950's, also witnessed lively debates about its nature and interpretation.¹ The arguments revolved around issues like: Is probability an objective feature of the phenomena under study, or merely a subjective judgment of the decision maker? How is probability related to frequency? If probability is an objective feature of reality, like heat or magnetism, then what scientific experiment could be devised to prove its existence and ascertain its value? If it is, on the other hand, a decision maker's purely subjective state of mind, then is there a way to judge its reasonableness or consistency with empirical evidence?

Classic works by Kolmogoroff (1950), Doob (1953) and Savage (1954) bypassed these philosophical issues by providing elegant mathematical formalisms of the concept of probability and related constructs. The phenomenal growth of modern probability theory and applications owes much to these works, which freed researchers from being bogged down with the hard conceptual issues of an earlier generation.

But setting foundational questions aside neither implies that these questions have been answered nor that their practical and conceptual implications magically disappear. In fact, we would argue that it is in management sciences, be it competitive strategy, finance, economics, or game theory, that

¹The classic works include Keynes (1921), Borel (1964), Knight (1921), Ramsey (1931), de Finetti((1937), (1989)), von Mises (1957), Reichenbach (1949), Savage (1954). Galavotti (2005) provides a comprehensive overview of the subject; see also Bernstein (1996) for a popular account of the notions of risk, uncertainty and probability. For the early subjectivist views of Ramsey and de Finetti, see Zabell (1991), Galavotti (1989), and Galavotti (2001).

these foundational questions about the meaning and interpretation probability have potentially the greatest significance.

To illustrate, consider some elementary issues in study of competitive strategy:

- *Over-optimism*: Do decision makers (firms, managers, investors) tend to be over-optimistic? But if probability judgments are purely subjective, then in what sense could they be wrong, or over-optimistic?
- *Objective vs. subjective probability*: Relatedly, are some probability judgments more ‘objective’ than others? Is there a sense in which decision makers can separate the objective from the subjective parts of their judgments?
- *Risk vs. uncertainty*: Should decision makers approach one-of-a-kind, highly uncertain decisions, like betting on the success of a new disruptive technology, in the same way as they approach routine, well-understood risks? More broadly, can one give formal meaning to Knight (1921)’s distinction between risk (roughly, events with known odds) and uncertainty (unknown, or unknowable odds)?
- *Belief formation, learning and testing*: How should firms translate their experiences into probability judgments to use in future decisions? And, is there a meaningful way to test these judgment against new evidence?

This is a sample of questions that surface, in different guises, in every major branch of the management sciences. An important role of a formal decision theoretic framework is to provide a systematic way to answer such questions.

The centerpiece of our survey is Savage (1954)’s theory of subjective probability, which remains to this day the foundation of the subject. We provide an account of its main assumptions and conclusions. Our emphasis is on the interpretation of the axioms, important attempts to extend the theory, possible critiques, and potential limitations.

2 Expected Utility Theory

2.1 Von Neumann-Morgenstern Representation

Before discussing Savage's framework,² it is useful to take a detour into von Neumann and Morgenstern (1947)'s classic representation theory when probabilities are objectively given. In their model, the primitives are a set of consequences: $C = \{\dots, x, y, z, \dots\}$, *e.g.*, monetary outcomes, and the set of probability distributions \mathcal{P} , or lotteries, on C with finite support. An individual has a preference \succsim over \mathcal{P} that is reflexive, complete, transitive and (suitably) continuous.³ The crucial ingredient in von Neumann and Morgenstern's theory is the *independence axiom*: for all lotteries p, q , and z and real number $\alpha \in (0, 1)$

$$p \succsim q \iff \alpha p + (1 - \alpha)z \succsim \alpha q + (1 - \alpha)z.$$

This is an additivity property of the preference: if two compound lotteries, like those on the RHS of the equivalence above, share a common component $(1 - \alpha)z$, then this component can be removed without affecting the ranking of the remaining, possibly non-common, components (p and q in this case).

Von Neumann and Morgenstern show that a preference satisfies the axioms if and only if there is a utility function $u : C \rightarrow \mathbb{R}$, unique up to positive affine transformation, such that:

$$p \succsim q \iff \int_C u(c) dp(c) \geq \int_C u(c) dq(c).^4 \quad (1)$$

In words, the decision maker ranks lotteries by applying the expected utility criterion with respect to the von Neumann and Morgenstern utility function u . When C is a convex set of real numbers, the utility function u will embed the decision maker's risk attitude.

²In addition to Savage (1954)'s classic work, Fishburn (1970) provides a textbook account while Kreps (1988) is an excellent, very readable introduction to the subject. Much of our terminology and notation below follows Machina and Schmeidler (1992).

³This is known as the Archimedean axiom. See the references above for precise statement.

⁴Given our assumption that \mathcal{P} consists of distributions with finite support, the integral here is, in fact, a sum. We use the integral notion to emphasize symmetry with similar expressions in other representation theorems.

Von Neumann and Morgenstern laid the foundation for the expected utility criterion which is central to all subsequent developments in decision theory. A major drawback, however, is that a decision maker in their model is somehow presented with probability distributions to make choices from. In most cases of interest, such as in games of strategy or business plans with even moderate degree of realism and complexity, decision makers are not offered the opportunity to choose between gambles with objectively known probabilities. In such situations the question is: under what conditions would decision makers' 'mental models' or 'representations' of their environment take the form of a probabilistic belief and the expected utility criterion? We turn to this next.

2.2 Savage's Framework

The key ingredients in Savage's theory are:

- States: $\Omega = \{\dots, \omega, \dots\}$;

In principle, a state of the world ω is a complete specification of every conceivable aspect of the decision problem at hand. In Savage's words, a state is "a description of the world leaving no relevant aspect undescribed." While this may be a useful conceptualization of states in a foundational work, Savage is, of course, aware of the need for a more parsimonious notion of state in practical problems.⁵

- Events: $\mathcal{E} = 2^\Omega = \{\dots, A, B, E, \dots\}$;

An event E is a set of states. Events will usually refer to information available to the decision maker, so the event E will stand for the piece of information that "the state belongs to E ." Note that the set of events is the power set 2^Ω , so there is no a priori restriction on what set of states can constitute an event. In Savage's framework, such restrictions ought to be viewed as part of the objective constraints facing the decision maker, rather than inherent in the framework itself. See Section 3.2.

⁵He refers to these as "small worlds," which are coarsenings of the underlying complete state space.

- Consequences: $C = \{\dots, x, y, z, \dots\}$;

A consequence x is a complete description of all that is relevant to the decision maker’s utility. This usually includes monetary rewards and other measures of material payoffs, but can also contain social and/or psychological factors (such as fairness, guilt, and the like).⁶

- Acts: $\mathcal{F} = \{\dots, f, g, \dots\}$.

An act is a finite-valued function that maps states to consequences. The restriction to finite values is for technical and expository convenience, and is not essential for the theory.

2.3 Savage’s Axioms

Savage’s goal was to derive a representation of a decision maker’s choice behavior in which uncertainty is represented by a probabilistic belief about the unknown states. This will, among other things, provide an interpretation and foundation of probability as the decision maker’s degree of belief expressed in his observed choice behavior. This behavior is formalized as a preference \succsim on \mathcal{F} (formally, a binary relation on $\mathcal{F} \times \mathcal{F}$). Although in concrete decision problems choice is limited to some exogenously given feasible set of acts $\mathcal{B} \subset \mathcal{F}$, the framework assumes a preference that is defined on all acts and is independent of the particular feasible set.⁷

A crucial methodological aspect of Savage’s framework is its focus on observable choices. Cognitive processes and other psychological aspects of decision making matter only to the extent that they have directly measurable implications on choice—at least in principle, possibly only under idealized or hypothetical experimental or choice settings. The process of decision

⁶Savage writes that consequences “might in general involve money, life, state of health, approval of friends, well-being of others, the will of God, or anything at all about which the person could possibly be concerned. Consequences might appropriately be called states of the person, as opposed to states of the world.”

⁷Other models of choice, such as the minimax regret criterion proposed by Savage (1951), allow dependence on the feasible set.

making—what steps a decision maker takes, or what heuristics he employs—has no formal meaning within the theory. The process is only relevant in so far as it has measurable behavioral consequences, and these are summarized by the preference \succsim in the sense that one can, in principle, offer the decision maker the choice between any two acts f and g and directly measure what choice he makes.

Seven axioms, numbered P1 through P7, characterize preferences that have an expected utility representation. Below we formally state and comment on the most controversial axioms, P1-P4.

Axiom P1 (Ordering) \succsim is complete, reflexive and transitive.

The key part of this axiom is completeness; it requires the decision maker to be able to rank any conceivable pair of acts. This entails an ability to conceive and rank all C -valued functions on the state space, no matter how complex these functions may be. Completeness also implies that a rational decision maker cannot say: “I am unable to choose between f and g because the evidence available to me is insufficient and/or ambiguous.” Savage’s response to these objections presumably would be that in any real choice a problem, a decision must ultimately be made, even when the evidence is scant or imperfect.

Completeness has been questioned by a number of authors. Bewley ((1986), (2002)), for example, argues that one should not expect completeness to hold when there is ambiguity about the probabilities. Shafer (1986) argued for a constructive interpretation of subjective probability, under which completeness need not hold.

Axiom P2 (Sure-Thing Principle–STP) For all events E and acts f, g, h and h' ,

$$\begin{aligned} \left[\begin{array}{l} f(\omega) \text{ if } \omega \in E \\ h(\omega) \text{ if } \omega \notin E \end{array} \right] \succsim & \left[\begin{array}{l} g(\omega) \text{ if } \omega \in E \\ h(\omega) \text{ if } \omega \notin E \end{array} \right] \\ \implies & \left[\begin{array}{l} f(\omega) \text{ if } \omega \in E \\ h'(\omega) \text{ if } \omega \notin E \end{array} \right] \succsim \left[\begin{array}{l} g(\omega) \text{ if } \omega \in E \\ h'(\omega) \text{ if } \omega \notin E \end{array} \right]. \end{aligned}$$

This is perhaps the most central, and controversial, of Savage's Axioms. It states that the preference between two acts with a common extension outside some event E does not depend on that common extension. In the formalism above, consequences are determined according to f and g on the event E , and a common act h outside E . The STP says that the preference remains unaltered if we replace h by some other common extension h' .

The STP cuts in many different ways and has been at the center of most of controversies in this area. It implies that the preference is separable across events, a necessary property for it to have an expected utility representation, and has an interpretation in terms of the consistency of dynamic choice. Two classic objections to the Savage framework based on violations of the STP are the classic papers by Allais (1953) and Ellsberg (1961).

For the next axiom, we need the following definition: An event E is *null* if any pair of acts which differ only on E are indifferent. Below, identify a consequence x with the constant act that yields x in every state.

Axiom P3 (Monotonicity) *For all outcomes x and y , non-null events E and acts g ,*

$$\begin{bmatrix} x & \text{if } \omega \in E \\ g(\omega) & \text{if } \omega \notin E \end{bmatrix} \succsim \begin{bmatrix} y & \text{if } \omega \in E \\ g(\omega) & \text{if } \omega \notin E \end{bmatrix} \iff x \succsim y.$$

Roughly, if a consequence x is preferred to another consequence y , then this is so regardless of the state in which these consequences obtain. Conceptually, this axiom amounts to separating states, which are the object of uncertainty, from consequences, which is what matters for payoffs.

Objections to this axiom appeared in Karni, Schmeidler, and Vind (1983), Schervish, Seidenfeld, and Kadane (1990), Karni (1993), among others.

Axiom P4 (Weak Comparative Probability) *For all events A, B , and outcomes $x^* \succ x$ and $y^* \succ y$.*

$$\begin{bmatrix} x^* & \text{if } \omega \in A \\ x & \text{if } \omega \notin A \end{bmatrix} \succsim \begin{bmatrix} x & \text{if } \omega \in B \\ x^* & \text{if } \omega \notin B \end{bmatrix} \implies \begin{bmatrix} y^* & \text{if } \omega \in A \\ y & \text{if } \omega \notin A \end{bmatrix} \succsim \begin{bmatrix} y & \text{if } \omega \in B \\ y^* & \text{if } \omega \notin B \end{bmatrix}.$$

Since $x^* \succ x$, the first preference ranking means that the decision maker prefers to bet on the event A rather than B . We would like to interpret this to mean that he therefore judges A to be more likely than B , an interpretation that is possible only if this likelihood judgments is independent of the specific consequences $x^* \succ x$. This is precisely what the axiom asserts.

Savage requires three more axioms, P5 (Nondegeneracy), P6 (Small Event Continuity), and P7 (Uniform Monotonicity).⁸ These axioms are needed primarily for technical reasons, and so they do not carry the substantive weight of the others.

2.4 Savage's Subjective Expected Utility Representation

Savage's Representation Theorem (Savage, 1954) *A preference \succsim satisfies P1-P7 if and only if there is a finitely additive probability measure⁹ P and a function $u : C \rightarrow \mathbb{R}$ such that for every pair of acts f and g :*

$$f \succsim g \iff \int_{\Omega} u(f(\omega)) dP \geq \int_{\Omega} u(g(\omega)) dP. \quad (2)$$

Moreover, P is unique and u is unique up to positive affine transformation.

In words, the theorem says that a decision maker who satisfies the axioms

- Reduces all uncertainty about the states to a subjective probability measure P that reflects his beliefs;
- Ranks consequences according to a utility function u that reflects his taste;
- Evaluates acts according to the expected utility criterion.

⁸P5 asserts the existence of two non-indifferent outcomes, and P6 implies that the subjective belief is atomless (in particular, Ω must be infinite; Gul (1992) provides an alternative model with a finite state space). P7 is necessary to handle infinite-valued acts.

⁹See footnote 16 for a brief discussion of the role of finite additivity.

Savage’s theorem implicitly *defines* probability as the degree of belief implied by the decision maker’s choice over uncertain prospects. Probability is not an objective property of the world, but the decision maker’s subjective assessment of the likelihood of various events, as expressed in his willingness to make bets on those events. In Savage’s framework it makes no sense to talk about beliefs being ‘right’ or ‘wrong,’ overly optimistic or pessimistic, or whether some probabilities may be more ‘objective’ than others.¹⁰

An important implication of the theory is the *separation of tastes from beliefs*. To illustrate this, consider two decision makers who may disagree on how to rank various acts. In principle, their disagreement may have its source in how they assess the likelihood of various events, in their preference over consequences, or some complex interaction between the two. For example, imagine two leaders who may disagree on how to deal with an emerging nuclear threat by a rogue regime. Their disagreement may be the result of different assessment of how likely the regime is to develop weapon-grade fuel, and/or how undesirable it is to have such regime armed with nuclear weapons. Suppose now that the two decision makers satisfy Savage’s axioms, with P_1, P_2 and u_1, u_2 denoting their subjective beliefs and utilities respectively. Then Savage’s theorem separates tastes from beliefs because it asserts that there can be no source of disagreement beyond differences between either P_1, P_2 and/or u_1, u_2 .¹¹

¹⁰Anscombe and Aumann (1963) offer a simpler derivation of subjective probability provided one is willing to assume the existence of objective lotteries. Specifically, their acts take values in the space of probability distribution on consequences, $\Delta(C)$, rather than ‘pure’ consequences C as in Savage, and the decision maker is assumed to rank objective lotteries according to expected utility criterion. The key assumption here is an appropriate version of the independence axiom.

¹¹In Savage’s theory, the separation between tastes and beliefs is a consequence of the fact that the representation identifies P, u , and the way in which they should be combined (the expected utility criterion). Machina and Schmeidler (1992) provides an alternative axiomatization— which relaxes the STP but strengthens P4—characterizing preferences that are *probabilistically sophisticated*. As in Savage, these are preferences in which the decision maker has a probabilistic belief P over the state space and who is indifferent between acts that induce the same distribution on consequences (under P). The difference is that lotteries over consequences are ranked according to a monotone functional; the expected utility criterion is but a special case of such functional. Machina and Schmeidler’s theory provides some of the weakest conditions known separating tastes from beliefs while not requiring expected utility.

What is the significance of separating tastes from beliefs? One may argue that both are aspects of the preference \succsim and both are subjective. Intuitively, however, tastes and beliefs are different animals. In Aumann (1987) words: “[U]tilities directly express tastes, which are inherently personal. It would be silly to talk about “impersonal tastes,” tastes that are “objective” or “unbiased”. But it is not at all silly to talk about unbiased probability estimates, and even to strive to achieve them. On the contrary, people are often criticized for wishful thinking—for letting their preferences color their judgement. One cannot sensibly ask for expert advice on what one’s tastes should be; but one may well ask for expert advice on probabilities.”

While Aumann’s reasoning is compelling, and undoubtedly shared by many (including the authors), it is important to understand that Savage’s theory has nothing to say about judging whether beliefs are reasonable or to test them against evidence. De Finetti points out that a subjective proposition, such as a subjective probability assessment, is one which “no experience can prove [...] right, or wrong; nor, in general, could any conceivable criterion give any objective sense to the distinction [...] between right and wrong.” (de Finetti (1989, p. 174)). The problem of developing criteria for testing beliefs, or to narrow the range of reasonable beliefs based on, say, criteria of simplicity, remains open.

3 Interpretation and Implications

3.1 Normative Theory and the Definition of ‘Rationality’

It is hard to argue that Savage’s theory and his representation accurately describe how people actually make choices. Classic examples by Allais (1953) and Ellsberg (1961) show that seemingly reasonable choices may violate the Savage axioms. A parallel development, pioneered by Tversky and Kahneman, draws on research in psychology to point out systematic deviations from Savage-style rationality. In a seminal paper, Tversky and Kahneman (1974) argue that “people rely on a limited number of heuristic principles which reduce the complex tasks of assessing probabilities and predicting values to

simpler judgmental operations. In general, these heuristics are quite useful, but sometimes they lead to severe and systematic errors.”¹² In general, the resulting behavior is inconsistent with Savage’s theory.

But if this theory is not descriptive of the way people actually make decisions, then what is its contribution? One answer is that the theory provides an operational *definition* of what we mean by rational behavior, namely as behavior consistent with the axioms, and thus has the expected utility representation in (2). This is a definition with far reaching consequences both in what it includes as well as what it excludes.

On the one hand, Savage’s theory may be seen as too permissive a framework for defining rationality. Any probability measure on the state space, no matter how absurd, qualifies as rational belief. Believers in Intelligent Design and members of the Flat Earth Society¹³ are rational, provided only that they are consistent with the representation. Savage is simply not in the business of passing normative judgments about what is and isn’t a reasonable belief, or how beliefs should be formed. Rather he seeks criteria for coherence of beliefs “to distinguish between coherent behavior and blunder, or demonstrable incoherence in the face of uncertainty” and it is best thought of as a tool “by which a person can police his own potential decisions for incoherency.” Savage (1967, p. 307).

On the other hand, Savage’s theory has powerful consequences in terms of what it rules out as irrational behavior, or as admissible models of such behavior. His framework, among other things: (1) erases any difference between risk and uncertainty; (2) expresses dynamic decision problems as constrained static problems; and (3) reduces differences of opinions to either unexplained difference in priors or to differences in information. These and other implications of Savage’s framework are discussed in later subsections. Here we illustrate the power of the framework in a particularly simple example.

Consider the following example which appears in Machina (1989): Mom has one indivisible treat to give to one of her two children, Abigail and Ben.

¹²See Kahneman, Slovic, and Tversky (1982) for a collection of important contributions to the literature on “heuristics and biases,” and Kahneman (2003) for a more recent account of the literature, especially in how it relates to economics and related fields.

¹³<http://www.alaska.net/~clund/edjublonskopf/Flatearthsociety.htm>.

Let a, b denote the acts of giving the treat to Abigail and Ben, respectively, and let m denote the act where Mom flips a fair coin and decides to give the treat to Abigail if the coin turns Heads and to Ben otherwise. Suppose that $a \sim b$, so Mom is indifferent between giving the treat to one child or the other. Expected utility implies that m should be indifferent to a and b , but it also seems entirely reasonable that $m \succ a \sim b$, the latter ranking indicating that Mom views determining who gets the treat based on a coin toss is fairer than, and therefore preferable to, arbitrarily assigning the treat to one child over the other. This, on the surface, seems like a perfectly rational behavior that contradicts Savage’s representation.¹⁴

The canonical answer to this and other anomalies within Savage’s framework is that instances of conflict between supposedly rational behavior and the rationality axioms must be accounted for as the result of a mis-specification of either the state space, the set of consequences, or the constraint set facing the decision maker. In Machina’s Mom example, the answer is easy: the set of consequences ought to take into account “fairness.” What matters for Mom and the kids is not just who gets the treat, but also whether the procedure is perceived as fair. The original, and apparently paradoxical, description of the problem missed an important payoff-relevant aspect: Mom may not be indifferent between different procedures for allocating the treat, and so the procedure should properly be part of the set of consequences.

Machina (1989, Section 6) provides examples and a critical discussion of the approach of change-the-consequences-so-anomalies-disappear presented in the last paragraph. He shows, among other examples, how enlarging the space of consequences can eliminate the famous “paradox” due to Allais (1953). In a related vein, Geanakoplos, Pearce, and Stacchetti (1989) provide examples and analysis of “psychological games.” These are strategic situations where payoffs may depend not just on the material consequences, but also on such considerations as fairness, guilt, vindictiveness, . . . etc.

The upshot of the above discussion is: (1) if the set of consequences (or

¹⁴Probabilities in this example are objective, in the sense of being exogenously given as part of the description of the problem. The example can, of course, be readily restated for subjective probability. The point of the example is to illustrate an intuitive contradiction with the reasonableness of the expected utility criterion aspect of Savage’s representation.

any other aspect of the model, such as the state space) is mis-specified, then there is no reason not to expect violations and anomalies; and (2) Savage’s framework provides a *methodology* that gives us guidance as to which models or explanations to pursue when we encounter anomalous behavior. In the above example, the framework suggests a mis-specification of the set of consequences as the source of the anomaly.

3.2 Feasibility Constraints and Complexity

Savage’s theory assumes a preference (and produces a representation) defined over all acts. In virtually every problem of interest, the decision maker chooses from a set of feasible acts, $\mathcal{B} \subsetneq \mathcal{F}$. In this case the representation identifies the optimal acts given \mathcal{B} as:

$$\operatorname{argmax}_{f \in \mathcal{B}} \int_{\Omega} u(f(\omega)) dP. \tag{3}$$

Although feasibility constraints do not formally appear in Savage’s representation, his theory has the implication that such constraints must be introduced as exogenous and objective elements of the decision problem. Subjectivity in Savage’s theory is limited to the decision maker’s taste over consequences, his beliefs over states, and the expected utility form he uses to combine the two. Objective feasibility constraints cannot, within Savage’s theory, influence the decision maker’s tastes or subjective judgments.

This has important implications. A common, and reasonable, complaint about Savage’s theory, *e.g.*, Shafer (1986), is that a decision framework should take into account the complexity of describing the states, of comparing acts, and so on. At one level, complexity, however defined, can be trivially accommodated within Savage’s theory: let $\mathcal{Z} \subset 2^{\mathcal{F}}$ denote the set of all feasible sets that are consistent with our favorite notion of ‘simplicity’ or computational feasibility, and apply (3) only to feasible sets $\mathcal{B} \in \mathcal{Z}$.

The problem with this approach is that in Savage’s framework, (3) makes sense only because we start with a preference over all acts and derive a representation which is independent of any cognitive or complexity-based constraints.¹⁵ One might reasonably expect cognitive or computational lim-

¹⁵Kopylov (2007) provides an extension of Savage’s theory to more general domains of

itations to constrain the decision maker’s preference. But to accommodate such constraints in Savage’s framework seems to require a decision maker who starts with a preference that violates them in the first place.¹⁶

3.3 Information, Dynamic Choice, and Dynamic Consistency

So what might be reasonable specifications of constraints that limit decision makers’ choices in practice? Obvious sources of constraints involve such things as wealth, technology, or regulations. Call these, for lack of a better term, *physical* constraints, so they can be more easily distinguished from *informational* constraints that reflect limited information. The simplest way to introduce limited information in Savage’s framework is in terms of a finite information partition $\Pi = \{\pi_1, \dots, \pi_n\}$ on the state space.¹⁷ At a state ω , the decision maker knows only that the state belongs to the partition element $\pi(\omega) \in \Pi$ that contains ω . Limited information, in the absence of other constraints of physical nature, can be viewed as constraining the decision maker’s choices to the subset of acts \mathcal{F}_Π that are measurable with respect to Π . An act $f \in \mathcal{F}_\Pi$ is a choice that is contingent on the information and can be evaluated in a completely standard way.

One may think of the above as representing dynamic choice. Savage’s original formulation is *timeless*, but it can be readily reinterpreted as a choice problem that takes place over time. Think of the problem as consisting of an

events that need not be algebras.

¹⁶A related issue is de Finetti’s and Savage’s insistence on assuming only finite, rather than countable, additivity. It is possible to obtain a representation like (2), as Arrow (1971) does, with countably additive probability, but at the cost of restricting attention to, say, a Borel σ -algebra of events, rather than the power set. Savage and de Finetti would object to this on methodological grounds: a decision framework should not blur the separation between structural assumptions about the choice setting from the feasibility constraints facing the decision maker in a particular choice problem. In Savage’s framework, imposing measure theoretic or topological structures, and the implied restrictions on the sets of events and acts, should be modeled as exogenous constraints on the set of feasible acts (e.g. Borel measurable acts), rather than as restrictions on the decision maker’s subjective beliefs. See Al-Najjar (2009) for discussion of this point.

¹⁷More sophisticated models represent information via a σ -algebra and, in dynamic settings, an information filtration.

ex ante stage, where acts are evaluated according to the representation (2), and an ex post stage where partial information about the state is revealed. Suppose this information takes the form that the state lies in some component π of the finite partition Π . Then the STP essentially amounts to saying that the ex ante preference \succ gives rise to a well-defined conditional preference \succ_{π} that is dynamically consistent in the sense that, roughly, its ranking of acts agrees with the ex ante ranking.¹⁸ Epstein and Le Breton (1993) show that dynamic consistency is essentially equivalent to a weakened version of the STP. In particular, models of ambiguity aversion (e.g. multiple prior models) cannot be dynamically consistent in any plausible sense (see section 3.4).¹⁹

The fact that Savage’s framework can so effortlessly accommodate limited information is one of its crucial advantages. Without this it would be hard to imagine how the development of games with incomplete information (Harsanyi (1967)), dynamic games of all types, and models of knowledge and interactive decision making (Aumann (1976), Aumann (1987)) could have proceeded.

What makes the formulation of dynamic choice within Savage’s model so attractive is its tractability: dynamic choice is nothing more than static choice subject to informational constraints. As noted by some authors (Kreps (1988) and Kreps (1998), for instance), this reduction to static choice seems to be missing intuitively important aspects of dynamic behavior we would care about as modelers in understanding games, political contests, or business strategy.

3.4 Risk vs. Uncertainty

The distinction between situation where probabilities are known and ones where there is insufficient information to form a probability judgment seems, at an intuitive level, meaningful. Most people think of coin tosses, dice

¹⁸Formally, the complement of π is null under \succ_{π} and for any two acts f and g that agree outside π , $f \succ g$ if and only if $f \succ_{\pi} g$.

¹⁹See also Karni and Schmeidler (1991). Wakker (1988) discusses another issue related to dynamic consistency, namely aversion to information that could result when the expected utility criterion is violated.

throws, or routine insurance losses as well-understood actuarial risks governed by processes with an objective probability. This is in contrast with events like “there will be a terrorist nuclear attack within the next 10 years,” or “my competitor will respond to my new product introduction with a price cut,” where there does not seem to be an obvious probability assignment. Knight (1921) and Keynes (1937) distinguish between situations of risk, where probabilities are well-understood, and problems involving uncertainty.²⁰ While this distinction seems, at some levels, quite intuitive, it has no meaning within Savage’s theory. Put differently, a decision maker in this theory reduces all uncertainties to risks, in the sense that he treats all uncertain prospects in the same way, namely by evaluating them using the expected utility criterion with respect to his subjective probability measure.

Bewley ((1986), (2002)), Schmeidler (1989), and Gilboa and Schmeidler (1989) extend the framework to allow for such distinction. Bewley relaxes the completeness part of P1 so the decision maker may be unable to rank some pairs of acts. He assumes that von Neumann-Morgenstern independence axiom holds over the set of acts on which his preference is complete. His main theorem represents such preferences as follows: There is a compact convex *set* of probability measure \mathcal{P}_B and a von Neumann-Morgenstern utility function u such that

$$f \succcurlyeq g \iff \int_{\Omega} u(f(\omega)) dP \geq \int_{\Omega} u(g(\omega)) dP, \forall P \in \mathcal{P}_B. \quad (4)$$

In other words, f is preferred to g if and only if there is unanimity in how they are ranked among all measures in \mathcal{P}_B . When the set \mathcal{P}_B is a sin-

²⁰ “By “uncertain” knowledge, let me explain, I do not mean merely to distinguish what is known for certain from what is only probable. The game of roulette is not subject, in this sense, to uncertainty; nor is the prospect of a Victory bond being drawn. Or, again, the expectation of life is only slightly uncertain. Even the weather is only moderately uncertain. The sense in which I am using the term is that in which the prospect of a European war is uncertain, or the price of copper and the rate of interest twenty years hence, or the obsolescence of a new invention, or the position of private wealth-owners in the social system in 1970. About these matters there is no scientific basis on which to form any calculable probability whatever. We simply do not know. Nevertheless, the necessity for action and for decision compels us as practical men to do our best to overlook this awkward fact and to behave exactly as we should if we had behind us a good Benthamite calculation of a series of prospective advantages and disadvantages, each multiplied by its appropriate probability, waiting to be summed.” Keynes (1937).

gleton, this reduces to standard subjective expected utility (which must, of course, be complete). A non-singleton \mathcal{P}_B represents uncertainty, or ambiguity, about the ‘true’ probability and will lead to incompleteness. The unanimity criterion in (4) is a robustness requirement under which the decision maker ranks acts only when such ranking is insensitive to how the ‘true’ prior varies within \mathcal{P}_B .

Gilboa and Schmeidler (1989) also develop a representation involving a convex set of priors.²¹ Rather than weakening completeness, Gilboa and Schmeidler (1989) weaken the independence axiom to obtain the following representation: there is a compact convex set of probability measure \mathcal{P}_{GS} and a von Neumann-Morgenstern utility function u such that:

$$f \succsim g \iff \min_{P \in \mathcal{P}_{GS}} \int_{\Omega} u(f(\omega)) dP \geq \min_{P \in \mathcal{P}_{GS}} \int_{\Omega} u(g(\omega)) dP. \quad (5)$$

That is, each act f is evaluated according to the worst measure $P_f \in \mathcal{P}_{GS}$ for that act. Again, this reduces to Savage-style representation if the set \mathcal{P}_{GS} is a singleton, in which case the preference must satisfy independence. However, when independence is violated (and their other axioms hold), the set \mathcal{P}_{GS} is non-singleton, and the resulting preference displays pessimism or caution. Gilboa-Schmeidler theory draws some of its motivation from the desire to reproduce choice patterns consistent with those appearing in the Ellsberg (1961) paradox. Their representation has been used to fit anomalies in asset pricing (starting with Dow and Werlang (1992)) and macroeconomics (Hansen (2007) for a recent example).²²

As attempts to capture ambiguity or uncertainty, Bewley ((1986), (2002))’s and Gilboa and Schmeidler (1989)’s representations suffer from significant problems. Bewley’s model cannot answer the obvious question: what should the decision maker do when he faces a choice between two acts f and g that are not ranked according to his criterion? Interesting cases of ambiguity arise

²¹In Schmeidler (1989)’s model, beliefs are represented as a capacity and the decision maker evaluates acts using the Choquet integral. These preferences overlap with those introduced in Gilboa and Schmeidler (1989).

²²See Al-Najjar and Weinstein (2009) for a critique of these efforts in the special issue of *Economics & Philosophy* on ambiguity aversion, and the replies to their paper in that same issue.

precisely when a decision maker has to choose between, say, two business plans about which he lacks precise priors. Presumably, the decision maker must ultimately make a choice, yet Bewley’s theory is silent on what this choice should be. It may be useful to think of Bewley’s model as corresponding to a different revealed preference experiment: while Savage envisions a setting where the decision maker is *forced* to make a choice when confronted with any feasible set \mathcal{B} , Bewley’s model is best thought of as recording the choice this decision maker would voluntarily make.

The problems with Gilboa and Schmeidler (1989)’s model are different but potentially more serious. It is well known, and can be easily demonstrated by example, that their model cannot admit any plausible form of dynamic consistency of choice. See Epstein and Le Breton (1993)’s paper, titled “Dynamically Consistent Beliefs must be Bayesian.” Arguments concerning the tension between their maxmin criterion (5) and consistent dynamic choice are discussed at length in Al-Najjar and Weinstein (2009).

It is interesting to note that neither Bewley nor the Gilboa-Schmeidler models appear to capture Knight (1921)’s intuition of what uncertainty is about: “[W]e must observe at the outset that when an individual instance only is at issue, there is no difference for conduct between a measurable risk and an unmeasurable uncertainty. The individual, as already observed, throws his estimate of the value of an opinion into the probability form [...] and “feels” toward it as toward any other probability situation.” This quote of Knight, and Keynes’ description of uncertainty in footnote (20), seems to imply that uncertainty is inherently about dynamic choice, while the Bewley and Gilboa-Schmeidler’s models produce uncertainty-sensitive behavior even in a single, static choice problem. This suggest that their representations capture a different type of behavior than what either Knight or Keynes had in mind.

3.5 Exchangeability, Objectivity, and Frequencies

As pointed out earlier, probability judgments in Savage’s theory are purely subjective; there is no sense in which some probabilities are more ‘objective’ than others. Intuitively, however, subjective probabilities should, somehow,

be grounded with objective facts, such as the frequencies of outcomes obtained in repetitions of the similar experiments. De Finetti’s notion of exchangeability and his famous representation make it possible to integrate a decision maker’s subjective judgment with objective frequencies. Our exposition here follows Kreps (1988) (who considers “de Finetti’s theorem as the fundamental theorem of (most) statistics.”)

Suppose that an experimental scientist or a statistician conducts (or passively observes) a sequence of observations in some set S (so the state space has the product structure $\Omega = S \times S \times \dots$). To pool information across experiments, the observer must assume, in one way or another, that the experiments are, in a sense, “similar.” De Finetti’s concept of exchangeability formalizes this intuition of similarity and its role as guide in decision making. Roughly, a decision maker subjectively views a set of experiments as *exchangeable* if he treats the indices interchangeably.²³

Different experiments will typically result in different outcomes due to the influence of a multitude of poorly understood or unmodeled factors. Nevertheless, a decision maker’s subjective judgment that the experiments are exchangeable amounts to believing that they are governed by the same underlying stochastic structure. De Finetti’s celebrated theorem says that a probability distribution P on Ω is exchangeable if and only if it has the parametric form:

$$P = \int_{\Theta} P^{\theta} d\mu(\theta). \tag{6}$$

Here the parameter set Θ indexes the set of all i.i.d. distributions P^{θ} with marginal θ , and μ is a probability distribution on Θ .

The decomposition (6) says that a decision maker subjectively views the experiment as exchangeable if and only if he believes the outcomes to be i.i.d. with parameter θ whose distribution is given by a subjective belief μ . Since the i.i.d. parameters may also be interpreted as limiting frequencies, the representation can be equivalently stated as: a subjective belief P is

²³Somewhat more formally, exchangeability means that the decision maker ranks as indifferent an act f and the act $f \circ \pi$ that pays f after a finite permutation of the coordinates π is applied. In particular, he considers an outcome s to be just as likely to appear in the i th as in the j th experiment.

exchangeable if and only if it can be expressed as distribution on limiting frequencies.

Chew and Sagi (2006) show that a form of event exchangeability (plus other weak axioms) implies probabilistic sophistication. De Castro and Al-Najjar (2009) isolate the implications of exchangeability assuming only that the preference is monotone, transitive and continuous, but otherwise incomplete and/or fail probabilistic sophistication. Their motivation is that de Finetti's representation simultaneously identifies the relevant parameters *and* Bayesian beliefs about them. On the other hand, the idea of similarity of experiments and the concept of parameters are meaningful even if expected utility is not assumed. De Castro and Al-Najjar develop a subjective version of the ergodic theorem which takes as primitive preferences, rather than probabilities (as in standard ergodic theory). They use this ergodic theorem to show a decomposition into i.i.d. parameters without necessarily assuming that the preference is expected utility.

4 Concluding Remarks

In his famous reply to calls to abandon a theory due to its anomalies and unintuitive implications, Hilbert said that “No one will drive us from the paradise which Cantor created for us.” Savage has created a ‘paradise’ where important ideas and, even entire fields, could flourish. It is hard to conceive of modern game theory, industrial organization, mechanism design, auction theory, models of incomplete information, bargaining theory, portfolio theory and option pricing without uncertainty and the use of probabilistic methods to reason about it. Savage has given us a broad license to reduce all uncertainty to risks that can be quantified using probability.

We argued above that one should set aside critiques of Savage that are based on misunderstanding or too narrow an interpretation of his framework. Often the failure is to recognize that Savage's axioms are less of substantive statement subject to refutations by experimental or psychological evidence, than an organizing methodology that guides and constrains our modeling of choice under uncertainty. By this standard, more than half century after its introduction, Savage's theory is an unqualified success. Yet there is a sense

that something was lost in the bargain, since we must give up the ability to talk about such things as the difference between risk and uncertainty, 'true' dynamic choice, or to have a meaningful theory of belief formation. These are promising direction for future research.

References

- AL-NAJJAR, N. I. (2009): “Decision Makers as Statisticians: Diversity, Ambiguity and Learning,” *Econometrica*, 77, 1339–1369.
- AL-NAJJAR, N. I., AND J. WEINSTEIN (2009): “The Ambiguity Aversion Literature: A Critical Assessment,” *Economics & Philosophy*, 25, 249–284, lead paper in a special issue on ambiguity aversion, with replies and a rejoinder.
- ALLAIS, M. (1953): “Le Comportement de l’Homme Rationnel devant le Risque: Critique des Postulats et Axiomes de l’Ecole Americaine,” *Econometrica*, 21(4), 503–546.
- ANSCOMBE, F., AND R. AUMANN (1963): “A Definition of Subjective Probability,” *The Annals of Mathematical Statistics*, 34(1), 199–205.
- ARROW, K. (1971): “Essays in the theory of risk-bearing,” *Amsterdam, London*.
- AUMANN, R. (1976): “Agreeing to disagree,” *The annals of statistics*, 4(6), 1236–1239.
- AUMANN, R. J. (1987): “Correlated Equilibrium as an Expression of Bayesian Rationality,” *Econometrica*, 55(1), 1–18.
- BERNSTEIN, P. L. (1996): *Against the Gods: The Remarkable Story of Risk*. John Wiley & Sons, New York.
- BEWLEY, T. (1986): “Knightian Decision Theory: Part I,” Cowles Foundation Discussion Paper no. 807.
- BEWLEY, T. (2002): “Knightian decision theory. Part I,” *Decisions in Economics and Finance*, 25(2), 79–110.
- BOREL, É. (1964): “Apropos of a Treatise on Probability,” in *Studies in Subjective Probability*, ed. by H. Kyburg, and H. Smokler, pp. 46–60. John Wiley and Sons.

- CHEW, S.H., C., AND J. SAGI (2006): “Event exchangeability: Probabilistic sophistication without continuity or monotonicity,” *Econometrica*, pp. 771–786.
- DE CASTRO, L., AND N. I. AL-NAJJAR (2009): “A Subjective Foundation of Objective Probability,” Northwestern University.
- DE FINETTI, B. (1937): “La prévision: ses lois logiques, ses sources subjectives,” *Annales de l’Institut Henri Poincaré*, 7, 1–68.
- (1989): “Probabilism,” *Erkenntnis*, 31(2), 169–223, (Originally published in 1931 as “*Probabilismo*” in Italian).
- DOOB, J. (1953): *Stochastic processes*. John Wiley & Sons.
- DOW, J., AND S. D. C. WERLANG (1992): “Uncertainty aversion, risk aversion, and the optimal choice of portfolio,” *Econometrica*, pp. 197–204.
- ELLSBERG, D. (1961): “Risk, Ambiguity, and the Savage Axioms,” *The Quarterly Journal of Economics*, 75(4), 643–669.
- EPSTEIN, L. G., AND M. LE BRETON (1993): “Dynamically consistent beliefs must be Bayesian,” *J. Econom. Theory*, 61(1), 1–22.
- FISHBURN, P. (1970): *Utility Theory for Decision Making*. John Wiley and Sons, New York.
- GALAVOTTI, M. (1989): “Anti-realism in the philosophy of probability: Bruno de Finetti’s subjectivism,” *Erkenntnis*, 31(2), 239–261.
- (2005): *Philosophical introduction to probability*. CSLI Publications Stanford, Calif.
- GALAVOTTI, M. C. (2001): “Subjectivism, Objectivism and Objectivity in Bruno De Finetti’s Bayesianism,” in *Foundations of Bayesianism*, ed. by D. Corfield, and J. Williamson, pp. 161–174. Kluwer Academic Publishers.
- GEANAKOPOLOS, J., D. PEARCE, AND E. STACCHETTI (1989): “Psychological Games and Sequential Rationality,” *Games and Economic Behavior*, 1, 60–79.

- GILBOA, I., AND D. SCHMEIDLER (1989): “Maxmin Expected Utility with Nonunique Prior,” *J. Math. Econom.*, 18(2), 141–153.
- GUL, F. (1992): “Savage’s theorem with a finite number of states,” *Journal of Economic Theory*, 57(1), 99 – 110.
- HANSEN, L. (2007): “Beliefs, doubts and learning: Valuing macroeconomic risk,” *American Economic Review*, 97(2), 1–30.
- HARSANYI, J. C. (1967): “Games with incomplete information played by “Bayesian” players. I. The basic model,” *Management Sci.*, 14, 159–182.
- KAHNEMAN, D. (2003): “Maps of Bounded Rationality: Psychology for Behavioral Economics,” *American Economic Review*, 93(5), 1449–75.
- KAHNEMAN, D., P. SLOVIC, AND A. TVERSKY (1982): *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge University Press.
- KARNI, E. (1993): “A definition of subjective probabilities with state-dependent preferences,” *Econometrica*, 61(1), 187–198.
- KARNI, E., AND D. SCHMEIDLER (1991): “Atemporal dynamic consistency and expected utility theory* 1,” *Journal of Economic Theory*, 54(2), 401–408.
- KARNI, E., D. SCHMEIDLER, AND K. VIND (1983): “On state dependent preferences and subjective probabilities,” *Econometrica*, 51(4), 1021–1031.
- KEYNES, J. (1921): *A treatise on probability*. Harper & Row.
- KEYNES, J. M. (1937): “The General Theory of Employment,” *The Quarterly Journal of Economics*, 51(2), 209–223.
- KNIGHT, F. (1921): *Risk, Uncertainty and Profit*. Harper, New York.
- KOLMOGOROFF, A. (1950): “Foundations of the theory of probability,” .
- KOPYLOV, I. (2007): “Subjective probabilities on “small” domains,” *Journal of Economic Theory*, 133(1), 236–265.
- KREPS, D. (1988): *Notes on the Theory of Choice*. Westview Press.

- KREPS, D. (1998): “Anticipated Utility and Dynamic Choice,” in *Frontiers of Research in Economic Theory: The Nancy L. Schwartz Memorial Lectures*, ed. by D. Jacobs, E. Kalai, and M. Kamien. Cambridge University Press.
- MACHINA, M. (1989): “Dynamic Consistency and Non-Expected Utility Models of Choice Under Uncertainty,” *Journal of Economic Literature*, 27(4), 1622–1668.
- MACHINA, M. J., AND D. SCHMEIDLER (1992): “A more robust definition of subjective probability,” *Econometrica*, 60(4), 745–780.
- RAMSEY, F. (1931): “Truth and Probability,” *The Foundations of Mathematics and Other Logical Essays*, pp. 156–198.
- REICHENBACH, H. (1949): *The theory of probability*. University of California Press.
- SAVAGE, L. J. (1951): “The Theory of Statistical Decision,” *Journal of the American Statistical Association*, 46(253), 55–67.
- (1954): *The foundations of statistics*. John Wiley & Sons Inc., New York.
- (1967): “Difficulties in the theory of personal probability,” *Philosophy of Science*, 34, 305–10.
- SCHERVISH, M., T. SEIDENFELD, AND J. KADANE (1990): “State-dependent utilities,” *Journal of the American Statistical Association*, pp. 840–847.
- SCHMEIDLER, D. (1989): “Subjective Probability and Expected Utility Without Additivity,” *Econometrica*, 57(3), 571–587.
- SHAFFER, G. (1986): “Savage revisited,” *Statist. Sci.*, 1(4), 463–501, With comments and a rejoinder by the author.
- TVERSKY, A., AND D. KAHNEMAN (1974): “Judgment under Uncertainty: Heuristics and Biases,” *Science*, 185(4157), 1124–1131.

VON MISES, R. (1957): *Probability, statistics and truth*. Allen and Unwin: Macmillan.

VON NEUMANN, J., AND O. MORGENSTERN (1947): *Theory of games and economic behavior*. Princeton university press Princeton, NJ.

WAKKER, P. (1988): “Nonexpected Utility as Aversion of Information,” *Journal of Behavioral Decision Making*, 1, 169–175.

ZABELL, S. (1991): “Ramsey, truth, and probability,” *Theoria*, 57, 211–238.