Nº 260
Marzo 2004

Documento de Trabajo

ISSN (edición impresa) 0716-7334
ISSN (edición electrónica) 0717-7593

Beyond Earthquakes: The New Directions of Expected Utility Theory.

Felipe Zurita

www.economia.puc.cl

Versión impresa ISSN: 0716-7334
Versión electrónica ISSN: 0717-7593

PONTIFICIA UNIVERSIDAD CATOLICA DE CHILE
INSTITUTO DE ECONOMIA

Oficina de Publicaciones
Casilla 76, Correo 17, Santiago
www.economia.puc.cl

BEYOND EARTHQUAKES: THE NEW DIRECTIONS OF EXPECTED UTILITY THEORY

Felipe Zurita*

Documento de Trabajo Nº 260

Santiago, Marzo 2004

---

* Instituto de Economía, Pontificia Universidad Católica de Chile. fzurita@faceapuc.cl

INDEX

Beyond Earthquakes:
The New Directions of Expected Utility Theory

Felipe Zurita[1]
Pontificia Universidad Católica de Chile

March 3, 2004

Abstract

Over the past two decades or so, an enormous amount of work has been done to improve the Expected Utility model. Two areas have attracted major attention: the possibility of describing unforeseen contingencies and the need to accommodate the kind of behavior referred to in Ellsberg's paradox. This essay surveys both.

[1] fzurita@faceapuc.cl I am very grateful to Burkhard Schipper, who was kind enough to patiently clarify my doubts regarding his work in a series of e-mails; to Federico Weinschelbaum for endless discussions; to Rodrigo Harrison and seminar participants at the Pontifical Catholic University of Chile for comments; and to Paolo Ghirardato for a long talk some years ago that brought this subject to my attention.

"But it does seem that the recent «incomplete market» studies may have left some of the most striking instances of incompleteness unstudied and in particular have given no indication how the Savage paradigm is to be deployed when we go beyond the weather and earthquakes."
Kenneth Arrow and Frank Hahn (1999)

1. Introduction

Expected Utility Theory is a cornerstone of economic analysis. It rationalizes individual decision making in ignorance, that is, when the decision maker does not know the consequences of each possible action. As situations of this kind are pervasive in economic life, it is not surprising that the theory has been so widely applied, to problems ranging from portfolio choice, gambling, insurance, investment, and education, to more intricate social phenomena such as coordination, delegation, or trading, to name just a few.

However, economists are finding a growing number of applications in which it seems incomplete and deficient, where observed behavior looks paradoxical through its lens. These paradoxes concern primarily more complex forms of ignorance. Consider, for instance, General Equilibrium Theory under uncertainty. Clearly, modern economies do not possess the type or the number of assets the model predicts, for it is extremely easy to find examples of untraded contingencies and in that sense, asset markets are incomplete. The study of incomplete market economies, however, has typically taken the incompleteness as a starting point, without explaining its source (with a few exceptions, notably Allen and Gale, 1989). Arrow and Hahn (1999) close their perceptive discussion of the internal consistency of the model with our opening quotation. A basic message of theirs is that economists have in some sense abused Expected Utility Theory (henceforth EUT), putting it to work at tasks for which it was not designed, such as analyzing the sophisticated forms of uncertainty that characterize asset market transactions. Over the past decade or so, much work has been done to produce a model that can handle these more complex forms of uncertainty, that is, that can go beyond earthquakes. This article surveys this work.

In particular, it focuses on two concepts that cannot be described in the language provided by Expected Utility Theory: unawareness and ambiguity.

The first refers to a situation in which the decision maker (henceforth DM) does not know the consequence of each available act and does not even imagine all possible consequences, and hence may eventually find himself in an unforeseen situation.

EUT's neglect of the higher level of ignorance posed by unawareness has important implications. For instance, a central result in Contract Theory asserts that two parties engaged in a long-term relationship will prefer to sign a complete contract over an incomplete one, meaning that all conflicts of interest that may arise in the relationship should be considered in the contract. This proposition is largely refuted by the evidence, as observed contracts are often vague and leave many contingencies undefined, which in turn often end up being renegotiated or legally disputed. Pioneered by Hart and Moore (1988), a branch of Contract Theory, Incomplete Contract Theory, has explored the implications that contracts could not be written on the variables of interest, say x, but on correlated variables, say y. Although this shortcut has allowed the study of many aspects of contracts, the foundations of this branch remain doubtful (Maskin and Tirole, 1999). The development of an Incomplete Contract Theory with solid foundations still awaits the development of a theory of behavior in conditions of unawareness. I describe in Section 3 some promising approaches on how to construct such a theory.

Speculation constitutes yet another example. The celebrated No Trade Theorem (Milgrom and Stokey, 1982) establishes that private information alone cannot be a motive for trade if there is common knowledge of rationality and all individuals have common priors. This latter assumption also involves their foreseeing a common set of contingencies; Heifetz, Meier and Schipper (2003) provide an example showing that this assumption is necessary for the No Trade Theorem. Thus, the development of an unforeseen contingencies theory may pave the way for a more intuitive and empirically relevant theory of speculation than we currently have.

In all these cases, extending EUT to accommodate the possibility of unawareness appears as a promising research area. This extension, however, has proved not to be simple, and the development of a tractable model is yet to come. Section 3 discusses the difficulties involved and promising new results as well.

The second strand of literature this article reviews is that of ambiguity, which in general terms investigates the behavioral consequences of non-additive representations of beliefs.

That the EU model relies on probabilities to represent beliefs is a fact. The implications of this fact, however, depend on the meaning attached to the mathematical concept of probability. For instance, to an objectivist observer, who feels that a probability is part of the description of an object or process, like weight or mass, EUT is essentially a normative theory whose positive value depends on its predictive ability. This analyst, however, may wish to distinguish situations in which the probabilities are known to the decision maker from those where they are not. For him, EUT only applies to the former class of situations, while the study of the latter, when the decision maker may have an *ambiguous* belief, requires a different tool.

A subjectivist, on the other hand, understands probabilities as an expression of the (personal) doubts the decision maker entertains in a particular decision problem. He uses the EU model because if the decision maker's behavior obeys certain properties, then he knows that he can represent it as the result of the maximization of an EU function, regardless of the decision maker's level of knowledge or ignorance. He would rely on the EU model to predict behavior as long as it is empirically satisfactory. It turns out that in some classes of problems, it is not. As a consequence, the theory needs to be amended. The choice problem occurs under *ambiguity* if the observed behavior can be better predicted by a non-additive belief.

How one interprets probability plays such an important role in determining the reasons behind attempts to generalize EUT's representation of beliefs that I have discussed the interpretation issue at some length. Keeping this multiplicity of meanings in mind also turns out to be crucial to understanding what EUT is – and what it is not. For this reason, I placed this discussion at the beginning, in Section 2.

Sections 3 and 4 discuss unawareness and ambiguity, respectively. The method consists of a general review of the theory, including its most commonly used applications and some recent advances. The review is not all-comprehensive, but focuses on identifying areas in which EUT's current use are problematic.

I should stress that this article is intended for a general audience within the profession. While there are good surveys available on these topics (like those of Dekel, Lipman and Rustichini, 1998b, on unawareness, and Ghirardato, 1993, on ambiguity) they tend to be too formal for most economists to fully grasp and apply. Instead, this article emphasizes interpretation, keeping the formal description to a minimum (although some is nonetheless necessary).

2. The Expected Utility Model

A choice problem occurs when a DM faces many mutually exclusive courses of action. Let A represent the set that contains all those acts. If the DM knew the consequence of each, he could evaluate those consequences directly, using, for example, function $V(a)$. Choice Theory assumes such an evaluation is possible and deduces the V function by looking at actual choices: the observed decision corresponds to the best evaluated among the available acts. Hence, if a^* was chosen when a was available, we infer that $V(a^*) \geq V(a)$ for all $a \in A$, or equivalently, that $a^* = \arg\max_{a \in A} V(a)$.

Uncertainty is a situation in which the DM does not know the consequence of each act. Although our usual image involves a DM who does not know what will happen in the future at the time when the decision is called for, the facts of which he is uncertain may or may not have already happened. Conceptually, whether or not the DM knows something has already happened or will happen is immaterial; all that matters is that the DM does not know about it. We can nevertheless identify an important dimension involving time: the moment at which the DM is ignorant and must decide (call it *ex ante*), another when the decision has been made and he may find out about something but not everything (call it interim), and a final moment when the DM learns what the consequences of his decision were (call it *ex post*). These three instants are intimately related in Decision Theory.

The EU model assumes that *ex ante* the DM is capable of listing all possible scenarios or states of the world. A state of the world is a description of all relevant aspects of reality that explain a given consequence. Knowing the state of the world is equivalent to knowing what the consequence of each act is or would have been, that is, a state of mind in which ignorance or uncertainty vanish. For instance, when the DM buys a lottery ticket, he does not know (he is uncertain) if he will win or not. He may list the possibilities: in the state of the world in which the winning number is his, he wins; in the other he does not; and it is irrelevant to this description whether the winning number has already been picked or not.

Say that Ω contains the list of all possible states of the world. Consider it a finite set, for reasons that will become clear later on. Observe that only one state is or will be "true" *ex post*, and all the rest will be declared to have looked plausible *ex ante* but proved "wrong" *ex post*. Label $c(\omega,a)$ the consequence of act a in state Ω and $u(c(\omega,a))$ the *ex post* evaluation of consequence $c(\omega,a)$ (which will obtain under act a if ω is the true state).

Observe that if only one state were considered possible, then the *ex ante* evaluation of the act may well be viewed as corresponding to the *ex post* evaluation of the consequence, that is $V(a)=u(c(\omega,a))$. This is in fact the usual convention in the timeless Choice Theory, where the *ex ante* and *ex post* instants are not distinguished. It is clear, however, that this cannot be so with uncertainty, as the evaluation of acts will vary with the states.

The EU model asks the DM in addition to hold numerically representable *ex ante* beliefs or degrees of trust in the truthfulness of each state. Moreover, these beliefs are assumed to be representable by a probability. Let $\pi(\omega)$ be the *ex ante* degree of confidence on ω being the true state.

The EU model holds that the DM behaves (i.e., chooses actions) as if maximizing a function such as this:

$$V(a)=E\left[u(c(\omega,a))\right]=\sum_{\omega\in\Omega}\pi(\omega)u(c(\omega,a)) \qquad (1)$$

that is, the function that evaluates acts (called utility function) corresponds to the expected value of the function that evaluates consequences (called the Bernoulli or felicity function).

The EU model is thus built on the following ideas: (i) The choice depends on an evaluation of the possible consequences (as opposed to principles, rules of thumb, etc., even though such ingredients may sometimes be considered part of the consequences); (ii) The list of possible consequences is complete, and we could in principle find a probability function that describes the DM *ex ante* beliefs affecting choices, and (iii) The evaluation of consequences and probabilities has the particular, linear form of Equation (1).

An axiomatization of EU is a set of axioms or assumptions about how DM behaves, which ensures (iii). There are many different axiomatizations of the EU model. Each, however, is built on a different idea of what a probability is, and hence in some sense constitutes a different theory. The most widely accepted axiomatizations of EU are also the polar cases in the dimension of interpretation: John von Neumann and Oskar Morgenstern's (1946) and Leonard Savage's (1957).

The main difference between them is their distinctive way of describing acts. While for von Neumann and Morgenstern an act is a probability distribution over consequences in itself, for Savage it is merely a map from states to consequences, deprived of any probability judgments in principle. For instance, if a DM is invited to bet on the outcome of a coin toss, von Neumann and Morgenstern would characterize the decision problem as composed of two mutually exclusive acts, described by:

$$\text{Accepting:} \quad \Pr(\text{winning } 1{,}000) = \Pr(\text{losing } 100) = \frac{1}{2}$$

$$\text{Rejecting:} \quad \Pr(\text{winning } 1{,}000) = \Pr(\text{losing } 100) = 0$$

while Savage would do it by:

Accepting: win 1,000 if heads, lose 100 otherwise
Rejecting: win/lose nothing in any event

Thus, in the treatment of von Neumann and Morgenstern, the probability is part of the description of the possibilities an individual faces, the courses of action open to him. Using this concept, the most natural interpretation of probability is of an objective nature: the probability exists outside the DM. A drawback is that since in any decision-making problem, the DM must know the alternatives from which he must choose, it follows that he must know the probabilities of each state. Depending on the application, this assumption may be too strong, and yet it is unavoidable if one wants to have a decision model at all.

Savage, on the other hand, characterizes acts as maps between states and consequences, with no probability judgments attached to them. In this approach, a DM understands an act as a list of possible consequences, one for each possible state of the world: $a : \omega \to c$. The most striking aspect of Savage's theory is that he manages to identify a set of behavioral assumptions that imply an expected utility representation of preferences, involving a probability function – which can be interpreted as the DM's *implied* (as opposed to declared or known) beliefs – and which, moreover, satisfies all properties of probabilities.

The significance of this fact is twofold. First, if one is willing to assume that an individual's behavior obeys these properties, then the existence of a probability function, to be interpreted as the DM's implicit beliefs, is granted. That is, the idea that an individual's beliefs are representable by a probability function does not involve a separate assumption. Specifically, it is not necessary to postulate anything about the DM's knowledge about empirical frequencies or the like.

On the other hand, Savage defines the notion of conditional preferences, to be interpreted as the preference ranking over acts after learning that an event has occurred. The representation of this conditional preference also takes the form of an expected utility, obtained from the unconditional one by an update of the probability that satisfies the definition of conditional probability. Hence, the implied beliefs satisfy Bayes' law, giving support to understand this as a dynamic theory in which inference from evidence is fully logical.

These polar axiomatizations, then, rest on a completely objective and a completely subjective view of probability, respectively. There are some intermediate views as well (e.g., Anscombe and Aumann 1963). Each of these axiomatizations provides a basis for a different theory of decision making under uncertainty, around a common mathematical model.

Since there are also many probability models, it is useful to state that the mathematical notion of probability referred to throughout this article is the mainstream formalization among probabilists, as per Kolmogorov (1950). For the sake of completeness, its basic structure is described briefly below.

2.1 The Probability Model

An event E is a subset of the set of states, Ω. E is perhaps more naturally defined as the set of states at which a given proposition p is true (e.g., $p \equiv$ "It will rain tomorrow"). Models that take as primitives the set of propositions are called *syntactic*, while those which take events as primitives are called *semantic.*[2] An algebra $\mathfrak{M}$ is a set of events, with the properties of being closed under complementation and union, that is:

(A1) $\varnothing, \Omega \in \mathfrak{M}$

(A2) $E \in \mathfrak{M} \Rightarrow E^c \in \mathfrak{M}$

(A3) $E, E' \in \mathfrak{M} \Rightarrow E \cup E' \in \mathfrak{M}$

An algebra is therefore the set of events describing the model's uncertain variables. From a syntactic perspective, the above properties involve the requirement that if one can talk about a given proposition, one should also be able to talk about its negation (A2); that there exists a proposition that is trivially true, or a tautology, and also its negation (A1); and that one can compound the available propositions to form new ones, with the logical connective "or" (A3) (and "and" by appealing to the de Morgan's laws, for $E \cap E' = \left(E^c \cup E'^c\right)^c$). Hence, if $p \equiv$ "it will rain tomorrow" and $q \equiv$ "the apple is rotten" are expressible in the model (i.e., $p, q \in \mathfrak{M}$), then "it will **not** rain tomorrow", "the apple is **not** rotten", "the apple is either rotten or not rotten", "it will not rain tomorrow but the apple is rotten", etc, also are.

A σ-algebra is an algebra that satisfies the following additional property:

(A3') If $\{E_i\}$ is an infinite collection of disjoint elements of $\mathfrak{M}$, then $E \equiv \bigcup_{i=1}^{\infty} E \in \mathfrak{M}$.

If Ω is finitely dimensional, clearly (3') is not different from (3); it only has bite over infinite dimensional state spaces.

---

[2] Aumann (1999a, b) offers a comprehensive but rather technical comparison of these models.

A **probability** π is a function that assigns to each event $E \in \mathfrak{M}$ a real number in the interval [0,1], $\pi : \mathfrak{M} \to [0,1]$, with the following properties:

($\pi 1$) $\quad \pi(\varnothing) = 0, \quad \pi(\Omega) = 1$ (normalization)

($\pi 2$) $\quad E \subseteq E' \Rightarrow \pi(E) \leq \pi(E') \quad \forall E, E' \in \mathfrak{M}$ (monotonicity)

($\pi 3$) $\quad \pi(E \cup E') = \pi(E) + \pi(E') - \pi(E \cap E')$ (finite additivity)

Depending on the interpretation, $\pi(E)$ can be a numerical judgment about the degree of truth in the proposition that E implicitly describes, or the degree of confidence in the occurrence of E, or the frequency at which E has been observed over some period of time. The structure imposed on π entails the requirements that tautologies (or the universe of possibilities or observations) receive the highest value (true or always) and their negation the lowest (false or never) (normalization); that the less demanding a proposition, the higher its degree of truth or frequency (monotonicity); and that if two propositions cannot be simultaneously true, then the proposition "either of them" should be assigned the sum of the degrees of truth of each individual proposition (additivity).

If one knew whether each proposition were true or false, one would know the "truth". There are many possible truths and the probability describes the degree of truth or confidence about each. Call each possible truth a state of the world – a full description of all relevant propositions – and label it ω. The union of all those states is a tautology if the list is so comprehensive (complete) that one can be certain that one of them must be true. Hence, this union must be Ω. Since each state is an event in its own right, describable as unions and intersections of other events, and since two states cannot be simultaneously true, the list of all states Ω can be used alternatively to describe $\mathfrak{M}$. Indeed, all events in $\mathfrak{M}$ can be written as unions of elements of Ω. This is why sometimes people (inappropriately) write $\pi : \Omega \to [0,1]$, because using this function, the probability of any event $E \in \mathfrak{M}$ can be computed as $\pi(E) = \sum_{\omega \in E} \pi(\omega)$ (where this last equality comes from additivity and the fact that $\{\omega\} \cap \{\omega'\} = \varnothing$ for any two different states).

Two remarks are in order. First, it should be clear by now that a probability is simply a real function with some properties and, as such, it may be useful to represent many different concepts. The most important of these are discussed in Section 2.2. EUT inherits this characteristic, for there are many interpretations that can be attached to the EU model as well. Second, the reference to propositions is not necessary to define events, and the semantic model can be constructed abstracting completely from any (explicit) syntactic model.

Paralleling conditions (A3) and (A3'), very often the latter property, additivity, is strengthened to require countable additivity, that is:

($\pi 3'$) If $E = \bigcup_{i=1}^{\infty} E_i$, where $\{E_i\}$ is a collection of countable many disjoint sets, then $\pi(E) = \sum_{i=1}^{\infty} E_i$.

Although the replacement of ($\pi 3$) by ($\pi 3'$) stirs some controversy among probabilists and philosophers, finite and countable additivity are equivalent in the finitely dimensional state spaces, which are the focus of this essay.

In the context of EUT, probabilities are used to represent beliefs, however acquired, more often than as a degree of truth. The probability model may account for dynamics of beliefs, by way of the concept of conditional probability. The probability of an event A conditional on an event B is defined by[3]:

$$\pi(A \mid B) = \frac{\pi(A \cap B)}{\pi(B)} \tag{2}$$

whenever $\pi(B) > 0$.

$\pi(A \mid B)$ is interpreted as the belief or degree of confidence in the occurrence of A after knowing that event B occurred. *Ex ante*, the DM thought that B could happen, but assigned a probability of $1 - \pi(B)$ to B not happening. After realizing B occurred, she must adjust her beliefs to account for the fact that now she knows B^c is impossible and will from now on assign it a probability of 0.

When applied to states that remain plausible, say $\omega, \omega' \in B$, the definition of conditional probability entails that their relative likelihoods do not change:

$$\frac{\pi(\{\omega\} \mid B)}{\pi(\{\omega'\} \mid B)} = \frac{\dfrac{\pi(\{\omega\} \cap B)}{\pi(B)}}{\dfrac{\pi(\{\omega'\} \cap B)}{\pi(B)}} = \frac{\pi(\{\omega\})}{\pi(\{\omega'\})} = \frac{\pi(\omega)}{\pi(\omega')} \tag{3}$$

Bayes' rule, on the other hand, connects $\pi(A \mid B)$ with $\pi(B \mid A)$ as follows:

[3] It might also be deduced from a set of axioms, much in the spirit of Savage' work. See, for example, Bernardo and Smith (1994), chapter 2.

$$\pi(A \mid B) = \frac{\pi(B \mid A)\pi(A)}{\sum_i \pi(B \mid A_i)\pi(A_i)} \tag{4}$$

where $\{A_i\}$ is a partition of Ω, that is, a collection of nonempty, pairwise disjoint sets, whose union is Ω. Observe that applying the definition of conditional probability in Equation (2), the rule reduces to the following statement:

$$\pi(B) = \sum_i \pi(B \mid A_i)\pi(A_i) \tag{5}$$

If $\{A_i\}$ is a list of mutually excluding events that the DM thinks of as possible *ex ante*, and if the DM observes *ex post* the occurrence of one such event, and if his beliefs about B change from $\pi(B)$ to some $\pi(B \mid A_i)$, then Bayes' rule is a consistency requirement between *ex ante* and *ex post* beliefs. Equation (5) asserts that the *ex ante* beliefs are the expected value of the *ex post* beliefs. If the DM were to believe something different *ex ante*, then he would use the information implicit on that belief to adjust some or all *ex post* beliefs.

I remarked earlier that in the work of von Neumann and Morgenstern probabilities are exogenous, and hence it is an additional assumption that they satisfy Bayes' rule. Whether it is an obvious, reasonable, or just plausible assumption is a matter, in part, of what interpretation is given to the notion of probability, as the next section discusses. In Savage's development, on the contrary, the satisfaction of Bayes' rule is the consequence of behavioral assumptions.

The axiomatic development of these theories and the proof of their respective representation theorems are widely available in graduate textbooks (e.g., Mas Colell et al.) and hence not given in this essay. What these sources do not emphasize (with the notable exception of Kreps' authoritative exposition, 1988), however, is the difference among available interpretations. An important source of differences stems from the disparate meanings attached to the term "probability," as surveyed below. A fuller discussion can be found in Hájek (2002).

2.2 Views on Probability

A. Objective Views

Objective views of probability hold that "probability" is a characteristic of the object under consideration, like weight, mass or temperature, and not of the observer. Presumably, the probability of an event can be learned through scientific examination, be it by empirical research or introspective reasoning.

The idea that probabilities are objective leads immediately to questions about the validity of EU theory as a positive theory. To be sustained, the theory needs to explain how these probabilities are "revealed" or "known" to decision makers.

The significance of this for EUT is that it severely limits its applicability to situations in which such probabilities can be computed. Several authors have argued against the possibility of applying this (or even any) mathematical model to explain behavior under uncertainty. In his General Theory, for example, Keynes distinguished between long-run and short-run expectation: the former referred to situations in which individuals can accumulate too little evidence to base judgments on it, thereby rendering probability assessments impossible. Knight, on the other hand, is credited for his distinction between risk (a situation in which the DM cannot predict the consequence but knows the probabilities) and uncertainty (where the latter knowledge is also absent), obviously holding an objective view.

Hence, from the perspective of EUT, objective views suffer from the common drawback that it cannot be taken as a universal, positive theory of human behavior, for in many, perhaps most, instances the relevant probabilities cannot be assumed to be known. This fact is of extraordinary importance and often overlooked.

Friedman's (1953) epistemological view, however, provides a partial reconciliation of objective views and universal applicability: the model must be judged by its predictions and not by its assumptions. The trouble is that from one model many predictions are drawn. If the researcher focuses on one prediction, then this may work. However, the discipline constantly seeks a general model that provides particular predictions. For instance, if the CAPM is taken to predict that expected returns are linearly related to their "betas", it may be of no consequence to assume that all investors "know" the joint distribution that characterizes asset returns. If it is taken to predict that all investors will hold a market portfolio, however, one is forced to look at the assumptions with more suspicion.

At any rate, universality may be the main practical reason behind why many economists declare their support for the subjectivist view. However, in many applications of EU theory in economics it is easy to recognize some, and often much, objectivism, a notable example being the Theory of Rational Expectations. A quick review of the most common objective interpretations follows.

i. The Classical View

Generally speaking the *classical* interpretation of probability holds that a DM should list all possibilities, consider each of them as "equally likely," and then deduce the probability of each event by "counting" the different ways that such event occur. The probability of an event is thus *defined* by the ratio "number of states consistent with the occurrence of the event/total number of states".

For instance, consider a gamble where a woman flips a coin once. If the outcome is heads, she pays \$2 to a bettor. If tails, she flips it again. If the outcome of the second trial is heads, she pays $\$2^2$ to the bettor. Otherwise she tosses it again, and so on. If, in the first n trials, the outcome was heads, the prize in case of tails in the (n+1)th trial is $\$2^{n+1}$. According to the classical interpretation, every gambler in this situation should consider the probability of winning exactly $\$2^n$ to be $\left(\frac{1}{2}\right)^n$. This conclusion is arrived at in the following way: "The number of outcomes in each trial is 2: the coin either comes up heads or tails. Hence, in each trial, the probability of heads is $\frac{1}{2}$. Therefore, after n trials, the possible outcomes take the form (x heads, n-x tails), and for each of these, there are $\binom{n}{x}$ sequences of heads-tails consistent with such outcome. However, the game ends the first time a tails appears. Therefore, to win exactly $\$2^n$, in the first (n-1) trial the outcome has to be heads and in the last one, tails. The probability of that is clearly $\left(\frac{1}{2}\right)^n$.

Possibly the earliest construction of EU is that of Bernoulli (1738), who proposed it as a solution to the Saint Petersburg paradox, one of the early discussions on the notion of value. The Saint Petersburg paradox used a gamble like the above to say that such a gamble has an infinite expected payoff, but clearly no one should be willing to pay even finite amounts such as, say, \$1,000,000. In effect, if $c(n)$ is the consequence of obtaining tails for the first time in trial n,

$$E\left[c(n)\right]=\sum_{n=0}^{\infty}\left(\frac{1}{2}\right)^n 2^n=\sum_{n=0}^{\infty}1 \tag{6}$$

which does not converge. Bernoulli, then, proposed the following "solution": if the gambler does not evaluate the state according to the prize he obtains, but rather according to a function $u(c(n))$, then he will be willing to pay a finite amount, as long as this function is concave. For instance, if $u(c)=\sqrt{c}$,

$$E\left[u(c(n))\right]=\sum_{n=0}^{\infty}\left(\frac{1}{2}\right)^n\sqrt{2^n}\approx 2.4, \tag{7}$$

a much more reasonable amount to pay for such a game. Of course, what Bernoulli constructed was the expected-utility function, giving simultaneously an argument for its use in decision theory (whether to pay an amount x for the right to play) and for what later became known as risk aversion, represented by the concavity of the Bernoulli function. Interestingly, then, EUT was born of a classical, objective view of probability.

There are many objections to the classical view of probability. Possibly the most devastating is the fact that any problem admits more than one description of the set of states. Since every state carries the same probability, different descriptions lead to different probabilities and hence to different preferences over acts. To give an obvious example, if the possibility of the coin landing on its edge is considered, then the probability of winning exactly $\$2^n$ should be $\left(\frac{1}{3}\right)^n$ instead of $\left(\frac{1}{2}\right)^n$. In the face of this, the theory does not uniquely predict what a DM would do.

ii. The Propensity View

The *propensity* view in some sense extends the notion of classical probability to allow for non-equal probability assignments, while retaining the idea that the probability is a characteristic of the object. Probability is then a deep, unknown parameter, which exists whether or not an observer is aware of its existence or value. Probability can be learned about by means of scientific investigation and the empiricist must experiment to discover its value. The deep parameter, however, can only be known precisely in an experiment that considers infinitely many repetitions.

Hence, a coin *propends* to exhibit a balanced number of heads and tails when tossed, in the same way a given light bulb *propends* to remain operational about 5,000 hours.

However, it is extremely unclear why these propensities must obey the properties usually attributed to probabilities (that is, normalization, monotonicity and additivity of either kind.) Each of these properties should in itself be the object of scientific inquiry. Yet, it is very hard to imagine a test for, say, normalization. Or more substantially, how to test whether the limit of a sequence of experiments is the purported distribution.

On the other hand, the definition entails a great deal of vagueness. Since all experiments must consider finitely many repetitions, then some notion of closeness must be adopted. These notions take the form of a set of conventions among researchers, which are often hard to defend – consider the magical number 30 as the minimum number of observations in regression analysis or the 5% considered the acceptable level of significance in statistical inference.

A more basic objection refers to the assumption that the experiments were made under "similar conditions." In many cases one might be forced to accept that if the conditions were literally the same, one should expect to see the same outcome. If, for instance, the coin was flipped two times, with the same force being applied in the same point, etc., it should land in the same place, in the same position. If this is true, then saying that 50% of the times it lands heads up merely says that the conditions were not the same in all coin flips.

The trouble is perhaps that often attached to the propensity view is the idea that random events exist in nature: the outcome of a coin flip is random, as is the duration of a light bulb or the price of a financial asset. In many instances, whether random variables exist in nature or only are artifacts to represent the observer's ignorance is of no consequence in research. Yet accepting the former means that random variables cannot be used to represent processes we know or feel are not random. The example of the previous paragraph suggested that the outcomes of coin flips are deterministic, although hard for humans to predict. Consider a different example: a child who is learning basic mathematics could certainly entertain doubts when faced with the question, How much is 322 plus 25? He may entertain many possibilities and may not be confident of his answer. Can his choice of answer be modeled by a EU maximization? As the result of the sum is not really random, an adherent to the first view would say no: random variables are for random events. A subjectivist, on the other hand, would be happy to do so.

In general, questions like "What is the probability that candidate X wins in the next poll?" are meaningless under this view. EUT, hence, would not be applicable to situations where probabilities cannot be defined.

iii. The Frequentist View

The *frequentist* view holds that probability refers to the frequency with which each particular event occurs, upon repeating an experiment under similar conditions a large number of times. It is in essence a mere, but convenient, description of the data.

As such, to the frequentist probability is not defined for one-shot experiments (or they are trivial: the probability of an outcome is either 0 or 1, depending on the observed outcome). As a consequence, this view supports a narrow application of EUT, as did the previous one.

iv. The Logical View

The *logical* view maintains that probability refers to the degree of truth of a proposition. Under this paradigm, probability is an extension of propositional logic, constructed over a syntactic model.

That such an extension can be satisfactorily made, and that it can be made by using probabilities, is a matter of controversy among logicians. But even assuming it can, it is not clear at all why individuals should base their behavior on these numerical values. For instance, a given proposition (event) may be considered to be more confirmed by evidence than another, and hence assigned a higher probability. Yet, EU requires the DM to behave differently when the ratio of these probabilities is, for example, 2, than when it is 3, although both cases conform equally to the notion that the first proposition is better confirmed by evidence than the second one. Put another way, the cardinality of probabilities is essential to decisions based on EU, although it is not necessarily so to represent degrees of confirmation.

B. The Subjective, Personalistic, or Bayesian, View

The *subjective* view of probability holds, on the other hand, that probability is a numerical representation of the degree of belief, or confidence, a DM has in a specific event occurring. The probability is thus a characteristic of the subject and not the object. It is an expression of his ignorance and doubts, in his present state of knowledge.

As such, it gains universality: every DM may hold beliefs and doubts about many events, be they random or not, if one insists on the existence of random events. It is not necessary to assume that the DM has experience with "similar" problems or that he knows relevant empirical frequencies. The subjectivist holds that, nonetheless, every DM may have an assessment, and that that assessment can be written in the form of a probability. This is not to say that experience, knowledge or information is useless to making an assessment: the theory would eventually explain belief-change in the presence of information, and in general, reconcile it with objectivist views in situations where a large number of experiments can be performed. It just does not limit itself to this kind of situation.

Thus, a DM may have beliefs about who is going to win the 2011 presidential election, despite the fact that the candidates have not even been officially nominated. Moreover, these beliefs are assumed to exist (even though we cannot think of a frequency distribution created from a large number of similar experiments), to differ from the "favorable cases/possible cases" formula (which would not only assign the same probability to each possible candidate, namely 1/10,000,000 in the case of Chile!), and to vary across people as well. Thus, besides its intuitive appeal and beautiful construction, the strong point of the subjective approach is its universal applicability to any decision problem. In the end, when faced with a decision problem, every DM makes a choice, so it is nice to have a theory that accounts for it or explains it.

The following informal summary of the meaning of "the probability of heads is $\frac{1}{2}$" using the different interpretations may be useful:

Table 1: The Meaning of the Phrase "The probability of heads is $\frac{1}{2}$."

View/School [Probability is] Meaning/Interpretation
Classical [A characteristic of the experiment of tossing a coin] Heads is one of the two possible outcomes of the experiment with "tossing a coin."
Propensity [A characteristic of coins] Coins propend to land heads up as much as they propend to land tails up when tossed.
Frequentist [Summary of evidence from many experiments] Experiments of tossing a coin a large number of times systematically result in balanced histograms for heads and tails.
Logical [Degree of truth] The degree of truth in the proposition "the coin will land heads" is the same as that of the analogous proposition regarding tails, before we have a chance to observe the outcome.
Subjective [A description of a person's beliefs] I see no evidence to favor one outcome over another, and thus would be willing to take any side on a bet on the coin landing heads.

C. Economists' Mainstream Interpretation

A paradox to EU as a positive theory is any systematic deviation from its predictions. Observe, however, that we can accommodate an EU representation of preferences in all surveyed views on probability, although with wider or narrower applicability to different decision problems, and that some phenomena may be contradictory to EU theory under a particular interpretation, while not under other interpretations.

Consider, for instance, the purchase of lottery tickets. Suppose that the gambler does it for the money, not the fun of gambling. Under a classical interpretation, such behavior is paradoxical if we further assume pervasive risk aversion, for such a transaction lowers the expected value and increases the risk of the individual's consumption bundle. Therefore, a EU maximizer would never do it, under a classical interpretation.

Under a subjective interpretation, on the other hand, there is no contradiction in assuming that all buyers believe they have higher chances of winning than other average buyers. The probability of winning is subjective and does not need to correspond to the ratio 1-over-the number of tickets sold. There is no conceptual problem in characterizing a situation in which all feel simultaneously luckier than the rest.

It should be stressed, however, that economists generally welcome the universality derived from Savage's development, but is reluctant to fully embrace the subjective view, mainly because there is the belief that taking subjectivism literally would render the theory almost empty: the "anything is possible'' critique.

In principle, there could be no connection between observed frequencies, however they may have been arranged, and beliefs. This is not to say that subjectivism holds that there is no link between beliefs and evidence, quite on the contrary. After all, Bayesianism obtained its name from Bayes' analysis of how rational people should interpret evidence and adjust beliefs accordingly. The point is that the learning procedure implied by Bayes' rule by itself does not restrict what an individual's beliefs should be at a certain point in time, even if we know what evidence he has been exposed to, without knowing his initial or *prior* beliefs, which need not be restricted by the data in any way.

Thus, popular applications of EU theory in all areas of economics impose either homogeneous beliefs – all individuals believe the same – or disparate information but common priors, so that everybody would believe the same if fed with the same information, the so-called Harsanyi doctrine. This is the case of rational expectation models in macroeconomics and finance (e.g., the CAPM, Grossman and Stiglitz, 1980)), and game-theoretical models of information (e.g., Spence, 1973, Riley, 1979, and Vickrey, 1961).[4]

Furthermore, when these models are analyzed empirically, the individuals' beliefs are taken to correspond to suitably defined empirical frequencies. This is the case, for instance, with the Efficient Market Hypothesis discussion.

The following was an actual classroom discussion in a finance course. The professor asked "What is the variance of tomorrow's price of IBM shares?" A student replied that next day's IBM share price did not have a variance, that while possibly every person was uncertain about the value, the beliefs of each could be represented by a probability function, which in turn could have an associated variance. Not knowing the value of the variance of an individual's beliefs amounts to saying that his beliefs are unknown to the observer; if the individual of interest is oneself, then not knowing the variance means that one is uncertain about one's own beliefs – but at this point, the question of what that variance is loses its original appeal. Clearly, the student adhered to a subjectivist view, while his professor to an objective, propensity view. What is interesting about this anecdote is that it illustrates that, although in many cases a question asked under one view can easily be rephrased to make sense in another, there are some research questions that seem important under one view but empty under another. This is a consequence of the fact that economists sometimes use a common model for radically different theories.

The professor's understanding of probability seems pervasive among economists. Incidentally, when asked to choose which of the five interpretations in Table 1 best represents their understanding of the phrase "the probability of heads is $\frac{1}{2}$", the five colleagues I asked chose the second one, and four doubted 2 and 3 were equivalent. Of course this does not prove anything, but illustrates the idea that if there is a mainstream interpretation, it is of an objective nature.

Apparently, then, economists welcomed the idea of applying EU pervasively to represent behavior under uncertainty, as it freed them from the narrow applicability sought by Keynes, for instance, but did not embrace the view about probability that supported such universality.

---

4 For a comprehensive discussion of the Common Prior Assumption in Economics, see Morris (1995). Gul (1998) and Aumann (1998) present opposite views on the role of subjectivism in economic theory.

Thus, although the interpretation and applicability of probability to economics was the subject of much discussion in the first half of the 20th century, by the second half the mathematical model had been accepted and was being widely applied (McGoun 1995). Since the debate over the meaning was never actually settled (although Savage's work, 1957, calmed down theorists), the new mathematical models received ambiguous interpretations. Empiricists, on the other hand, embraced propensity views, even when the more natural interpretations of the models might be subjective, because of the need to have ready counterparts to the individual's beliefs in the data.

2.3 EU's underlying knowledge structure

Regardless of the interpretation given to "probability," the assumption that a DM's beliefs can be represented by a probability function implies that the DM's knowledge at each point in time obeys a particular logical structure, namely, a Kripke-S5 system. A Kripke-S5 is a system of knowledge that satisfies what are generally considered highly demanding but desirable logical principles. Under objectivist interpretations, such requirements are fine, for probabilities should obey logic. Under subjective interpretations, they constitute assumptions over reasoning that preclude what could be labeled as "reasoning mistakes." This section explains these ideas in some detail. It should be stressed that the contents of this section are widely available elsewhere (see, for instance, Osborne and Rubinstein's textbook, 1994); they are included here because the concepts developed in this section will be used extensively in the next.

I have said that Ω contains a list of all possibilities that *ex ante* the DM envisions for a given future date (or all the hypotheses he considers could be true). Only one state will obtain (or is true). The DM may not learn the truth under all circumstances, however. This may be represented by means of a *possibility correspondence*,

$$\begin{aligned} P &: \Omega \to 2^{\Omega} \\ &\omega \to P(\omega) \end{aligned} \tag{8}$$

indicating at each state what states the DM considers *ex post* to be possible.

In the particular case in which the individual always discovers the true state, we have $P(\omega) = \{\omega\}$ for all $\omega \in \Omega$. In general, however, it is possible that the DM is not completely sure about the occurrence of a state and entertains more than one possibility.

It is standard in economic analysis to impose on P the requirement that the conclusions be "logical." This means that no state exists in which the DM is sure that the true state has not occurred, because it would be impossible to see conclusive evidence against the true:

(P1) $\forall \omega \in \Omega, \quad \omega \in P(\omega)$

If the DM does not see conclusive evidence against ω' when ω is the true state, then he should not see conclusive evidence against ω when ω' is the true state, for if he did, at ω' he could reject ω as a possibility, and knowing that, he could reject ω' at ω on the grounds that he cannot reject ω and hence he cannot be at ω'.

(P2) $\quad \forall \omega' \in P(\omega) \Rightarrow P(\omega') = P(\omega)$

Lemma $\quad P$ satisfies P1 and P2 if and only if it induces a partition on Ω.

Proof: See, for instance, Osborne and Rubinstein (1994), p. 68.

A possibility correspondence that satisfies these properties is, consequently, called partitional. An interesting feature of partitional possibility correspondences is that they imply a fully logical knowledge structure on the part of the DM.

Say that the DM *knows* that an event E obtained if E is true in every state he considers possible. Consequently, from P we can define the knowledge operator:

$$K(E) = \{\omega \in \Omega \mid P(\omega) \subseteq E\} \tag{9}$$

which indicates, for each event E, in what states the DM knows that E is true. Intuitively, since set inclusion is defined as implication of belonging (i.e., $\langle P(\omega) \subseteq E \rangle \Leftrightarrow \langle \omega \in P(\omega) \Rightarrow \omega \in E \rangle$), knowledge of $P(\omega)$ involves knowledge of all events it implies. Similarly, $\neg K$ indicates that the individual does not know whether E is true or not (which should not be confused with $K(\neg E)$, knowing that $\neg E$ is true.)

A system of knowledge (Ω, K) is a Kripke-S5 if it satisfies the following properties:

(K1) $\quad K(\Omega) = \Omega$ (necessitation)

(K2) $\quad K(E) \cap K(E') = K(E \cap E')$

(K3) $\quad K(E) \subseteq E$

(K4) $\quad K(K(E)) = E$ (positive introspection)

(K5) $\quad \neg K(E) = K(\neg K(E))$ (negative introspection)

K1 asserts that in all states, the DM knows the tautologies. K2 that the states for which he knows events E and E' are true are those in which he knows both events. K3 says that the states in which E is known do not include states in which it is false. These first three axioms are satisfied by many knowledge functions; the next two, in turn, only by partitional structures.

K4 requires the DM to be aware of his knowledge: if the DM knows E, he must know that he knows it. K5 further requires that if he does not know something, he knows that he does not know it.

Theorem: K satisfies K1-K5 if and only if P is partitional.
Proof: See, for instance, Osborne and Rubinstein (1994), p. 70.

These axioms amount to a strong requirement on the deductive capabilities of the DM. Consider the following example (taken from Dekel, Lipman and Rustichini, 1998a, referred to from now on as DLR, who in turn adapted it from Geanakoplos, 1989):

“While Watson never reported it, Sherlock Holmes once noted an even more curious incident, that of a dog that barked and the cat that howled in the night. When Watson objected that the dog did not bark and the cat did not howl, Holmes replied ‘that is the curious incident to which I refer.’ Holmes knew that this meant that no one, neither man nor dog, had intruded on the premises the previous night. For had a man intruded, the dog would have barked. Had a dog intruded, the cat would have howled. Hence, the lack of either of these two signals means that there could not have been a human intruder or canine intruder.”

Following DLR, say that a is the state in which there is a human intruder, b the state in which there is a canine intruder, and c the state in which there is no intruder at all. Watson’s set of states is therefore $\Omega=\{a,b,c\}$, and his possibility correspondence:

ω	$P_W(\omega)$
a	$\{a\}$
b	$\{b\}$
c	$\{a,b,c\}$

Table 2: Watson’s Possibility Correspondence

meaning that when there is a human intruder or a canine intruder, he is aware of its presence, but when there is no intruder he considers possible all three states, because he fails to realize that had a man or a dog intruded, he would know it, and therefore, as he does not know it, the state must be $\{c\}$. In contrast, the logically-sophisticated Holmes has the possibility correspondence P_H:

ω	$P_H(\omega)$
a	$\{a\}$
b	$\{b\}$
c	$\{c\}$

Table 3: Holmes’ Possibility Correspondence

To be sure, Watson's knowledge system is not a Kripke S5. In particular, his knowledge system fails to satisfy property K5, negative introspection, which can be appreciated by comparing the last two columns of Table 4. For instance, in state c he knows that he does not know $\{a\}$, but he does not know this fact: $K\left(\neg K\left(\{a\}\right)\right) \neq \neg K\left(\{a\}\right)$.

E	$P(E)$	$K(E)$	$K(K(E))$	$\neg K(E)$	$K(\neg K(E))$
$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$	Ω	Ω
$\{a\}$	$\{a\}$	$\{a\}$	$\{a\}$	$\{b,c\}$	$\{b\}$
$\{b\}$	$\{b\}$	$\{b\}$	$\{b\}$	$\{a,c\}$	$\{a\}$
$\{c\}$	$\{a,b,c\}$	$\varnothing$	$\varnothing$	Ω	Ω
$\{a,b\}$	$-$	$\{a,b\}$	$\{a,b\}$	$\{c\}$	$\varnothing$
$\{a,c\}$	$-$	$\{a\}$	$\{a\}$	$\{b,c\}$	$\{b\}$
$\{b,c\}$	$-$	$\{b\}$	$\{b\}$	$\{a,c\}$	$\{a\}$
Ω	$-$	Ω	Ω	$\varnothing$	$\varnothing$

Table 4: Watson's Knowledge Structure

An important feature of a non-partitional possibility correspondence is that the probabilities they generate fail to satisfy Bayes' rule in the following sense: if $\pi(E \mid \omega)$ is the degree of confidence on E when the true state is ω, then the analog to Equation (5) is

$$\pi(E) = \sum_{\omega \in \Omega} \pi(E \mid \omega)\pi(\omega) \tag{10}$$

However, $\pi(E \mid \omega) = \pi(E \mid P(\omega))$. Since the possibility correspondence is not partitional, then in general we have:

$$\pi(E) \neq \sum_{\omega \in \Omega} \pi(E \mid P(\omega))\pi(\omega) \tag{11}$$

In the example, if Watson regarded the three possibilities as equally likely, his prior beliefs in each event would be given by:

E	$\pi_W(E)$
$\varnothing$	0
$\{a\}$	$\frac{1}{3}$
$\{b\}$	$\frac{1}{3}$
$\{c\}$	$\frac{1}{3}$
$\{a,b\}$	$\frac{2}{3}$
$\{a,c\}$	$\frac{2}{3}$
$\{b,c\}$	$\frac{2}{3}$
Ω	1

Table 5: Watson's Beliefs

It can readily be checked that this function satisfies properties $\pi1-\pi3$. However, the probability of there being an intruder (event $\{a,b\}$) does not coincide with the sum of the corresponding conditional probabilities as Bayes' rule asserts:

$$\begin{aligned}\pi(\{a,b\}) = \frac{2}{3} &\neq \sum_{\omega=a,b,c} \pi(\{a,b\} \mid \omega)\pi(\{\omega\}) \\ &= \sum_{\omega=a,b,c} \pi(\{a,b\} \mid P(\omega))\pi(\{\omega\}) \\ &= 1*\frac{1}{3}+1*\frac{1}{3}+\frac{2}{3}*\frac{1}{3}=\frac{8}{9}\end{aligned} \tag{12}$$

The problem stems from the fact that when c is true, Watson does not realize it, and considers everything as possible. His conditional belief should be $\pi(\{a,b\} \mid \{c\})$ rather than $\pi(\{a,b\} \mid \{a,b,c\})$. It is clear, on the other hand, that Holmes' possibility correspondence is partitional, his knowledge function satisfies Kripke's S5 conditions, and his beliefs satisfy Bayes' rule.

Therefore, when we invoke an EU representation of behavior and impose (in the von Neumann Morgenstern case) or deduce (in the Savage case) the satisfaction of Bayes' rule for conditional beliefs, we are ruling out incomplete logical deductions like Watson's.

The next two sections discuss aspects of the EU hypothesis that seem problematic and remain in search of an answer.

3. Unawareness

Thus far, we have characterized an EU maximizer as a fully logically-consistent individual – either because his subjective beliefs have that property or because he knows the objective probability function that logically satisfies that property. This section holds on to that idea, but explores the possibility that his knowledge of the problem at hand is incomplete, in the sense that some possibilities may not cross his mind at the time the decision is called for. This does not mean that the likelihood of this event was deemed negligible, but rather the more extreme situation in which the DM does not foresee a possibility or is *unaware* of it.

Under a subjectivist view, the idea of unawareness complements the notion of ignorance quite naturally. If an individual ignores what will happen in the future, we say that he is uncertain. If he does not know everything that could happen, we say that he is unaware of some possibilities. Note that we will say that anything the DM imagines as possible is indeed possible, while anything that is *ex post* observed to have happened but was not considered *ex ante*, was possible – an unforeseen contingency.

In introductory examples of EU theory this case does not arise: the DM is uncertain of the weather the following day and within the example it is obvious that rain/not rain exhaust all the possibilities. Or the DM is uncertain about whether an earthquake will occur before year's end. Or, at most, he is uncertain about the level of consumption he will achieve if he invest in a particular portfolio, and the possibilities are well covered by the positive real numbers. In all these examples, it would appear as logically impossible that a state from outside the specified state space materializes.

But it is not. All those descriptions reflect the way the DM sees the problem at hand with the information and resources available at the time of decision. However, his *ex post* evaluation of the act may as well be based on variables not considered *ex ante*. For instance, when purchasing an umbrella, a woman did not foresee that by choosing a fancy color she would be chased by confused bees. Her *ex ante* description of the problem did not include the bees that conformed her *ex post* account of the consequence of her decision.

Moreover, some decision problems are better described by an incomplete set of consequences. Consider the case of scientific research, where the outcome is not only uncertain, but often unimaginable. Similarly, consider the case of a writer who is planning to write a novel. She does not know what the finished novel will look like. So let's say she lists all possibilities: in her mind this must be all the books in the world, limited perhaps to her language, written, to be written, never to be written. This is logically plausible, as the writer knows all the letters and can certainly imagine all combinations of the characters,[5] so in principle she could imagine all possible books and the process of writing would involve choosing from among those possibilities. The same would be true, of course, of the outcome of scientific research reports in the form of papers. Or more generally, of any ideas amendable of written communication, including of course, music.

---

[5] Reminiscent of the problem presented in Borges' short story "The Library of Babel" (in "The Garden of Forking Paths", 1941).

This is clearly not possible, perhaps because of the brain's or cognition's physical restrictions. In the case of the writer, she probably cannot hold even one whole book in her mind at the same time. Observe, then, that the possibility of unforeseen contingencies is latent even in finitely dimensional state spaces (as in the book example, if we take the number of pages of every book to be bounded.)

In these examples, learning is not of the shrinking nature presented by Bayes' rule, where the DM uses the evidence to discard possibilities, but of an expanding nature: evidence may not help the individual discard states initially considered possible, which then do not materialize, but instead reveal other unforeseen events that actually happened or could happen in the future.

The reader may object to this as a criticism of EUT by pointing out that predictive, rather than descriptive accuracy is its purpose. Indeed, the fact that EUT does not describe the process of writing cannot be considered a major objection if it still predicts well. And yet, should not the model that predicts the best available contract between two parties anticipating future conflicting interest include a complete description of what each should do in every state of nature? As emphasized earlier, some EUT predictions do indeed refer to the description of contracts, laws, and market organization, issues in which the *ex ante* indescribability of future contingencies is crucial. One must be aware of an event in order to have a belief about it and make a decision associated with it.

In Subsection 2.1 I discuss some difficulties involved in the extension of the EU model to incorporate unawareness and some possibly fruitful applications in Subsection 2.2.

2.1 The Description of Unawareness

When looking at examples such as that of Watson, some authors (e.g., Geanakoplos 1989) conjectured that it would be possible to extend the state-space model to account for unawareness, by departing from partitional possibility correspondences. Yet, the task has proven not to be so simple.

Certainly, in state c, when the dog does not bark, Watson seems to be unaware of this. If $P(\omega)$ is reinterpreted as the states the DM cannot discard (either because he thinks they are possible or because he is unaware of them), then $P(c)$ would include the possibilities of the dog barking and the cat howling, as in Table 2.

Paralleling the definition of the knowledge operation, one may wish to construct an unawareness operator $U(E)$, indicating at what states the DM is unaware of an event E.

It seems natural to require that if the DM is unaware of an event E, then he should not know E nor know that he does not know it, for the latter would imply that the event is in his mind. Thus, a possible definition (although some authors have suggested the need to strengthen it) of unawareness is given by:

$$U(E) \equiv \neg K(E) \cap \neg K(\neg K(E)) \tag{13}$$

We can check in Table (6) (continued from Table 4) that at c Watson seems to be unaware of $\{a\}$: not only at that state does he not know of the event $\{a\}$ ($c \notin K(\{a\}) = \{a\}$), but also he does not know that he does not know it ($c \notin K(\neg K(\{a\})) = \{b\}$), that is, he is unaware of his lack of knowledge.

E	$\neg K(E)$	$\neg K(\neg K(E))$	$U(E)$
$\varnothing$	Ω	$\varnothing$	$\varnothing$
$\{a\}$	$\{b,c\}$	$\{a,c\}$	$\{c\}$
$\{b\}$	$\{a,c\}$	$\{b,c\}$	$\{c\}$
$\{c\}$	Ω	$\varnothing$	$\varnothing$
$\{a,b\}$	$\{c\}$	Ω	$\{c\}$
$\{a,c\}$	$\{b,c\}$	$\{a,c\}$	$\{c\}$
$\{b,c\}$	$\{a,c\}$	$\{b,c\}$	$\{c\}$
Ω	$\varnothing$	Ω	$\varnothing$

Table 6: Watson's Unawareness Structure

DLR show, however, that this example has an awkward feature: while $U(\{a\}) = \{c\}$, that is, at c Watson is unaware of $\{a\}$, at the same time $U(U(\{a\})) = U(\{c\}) = \varnothing$, that is, Watson is aware (i.e. not unaware) of the fact that he is unaware of $\{a\}$. Moreover, this feature is not specific to this example, but common to standard state-space models.

This is a devastating result that militates against the idea of amending the model, as suggested by Geanakoplos (1989). There are other ways, nonetheless. Below we will review two (currently unpublished) propositions. Their common element is to stress the need of incorporating into the state the description of the DM's epistemic status.

Rather than proposing a new model, Ely (1998) presents a reconciling interpretation. He argues that the breakdown of negative introspection is not intended to reflect logical inconsistencies in the reasoning of the DM, but rather his failure to capture "the whole picture" (which the modeler sees), i.e., his lack of understanding of the way his perception works, which also happens to be the source of his unawareness.

Ely conjectures that this can be captured by a model based on a collection of partitions of Ω, $\{\Pi_\omega\}$, one for each state, representing the DM's understanding of the situation at each state. The modeler knows that at state ω the thinking of the DM will be characterized by $\Pi_\omega(\omega)$, although at that state, the DM thinks that at state ω' he would have thought $\Pi_\omega(\omega')$. Since at each ω his possibility correspondence is a partition, the DM's model is internally consistent. Thus, Watson's case could be modeled in the following fashion:

ω	$P_a^W(\omega)$	$P_b^W(\omega)$	$P_c^W(\omega)$
a	$\{a\}$	$\{a,c\}$	$\{a,b,c\}$
b	$\{b,c\}$	$\{b\}$	$\{a,b,c\}$
c	$\{b,c\}$	$\{a,c\}$	$\{a,b,c\}$

Table 6: Watson's Possibility Correspondences in Ely's View

Observe that while the modeler "sees" the main diagonal, reconstructing Table 2 and recognizing Watson's violation of negative introspection, Watson does not know what he would have thought in other states, and is unaware of some possibilities. At state c he thinks everything is possible and holds the logically consistent (but incorrect) belief that he would have thought the same had state a occurred.

Under this approach, the criticism that DLR raises over the unawareness operator of Equation (13) is no longer valid: when applied to the modeler's understanding of the situation, it is fine for him to say that at c Watson is unaware of $\{a\}$, and the modeler is aware of this fact, but not Watson. Observe that at c, $K_c(\{a\}) = \varnothing$, and also $K_c(\{b,c\}) = \varnothing$, so in no state does he know that there was a human intruder or its negation.

Schipper (2002), on the other hand, proposes rewriting the model, which not only allows for a form of unawareness immune to the DLR critique, but also has the advantage of being capable of incorporating multiple individuals, and speaking of "interactive unawareness." This is certainly a must for the theory to be useful for analyzing social problems. Nonetheless, I will focus on individual decision making here.

Schipper introduces a crucial modification to the state-space model, namely, a redefinition of the concept of state, in two dimensions. In the first place, he includes in the state the description of the epistemic status of the individual (which obviously leads to an enlargement of the state space.) There are three possible epistemic statuses regarding a fact or proposition: the DM may be **unaware** of it; he may be aware of it but not necessarily know its "value" (call it **know whether**), or lastly he may know its value (call it **know that**). In principle, then, we could have as many states as combinations of these three status with the "old" (or external) states. In the second place, he allows the state to be an incomplete description of the (external variables of the) world. Indeed, while for Savage a state is a complete description of all relevant variables in a decision problem, for Schipper it will be a complete description from the DM's point of view and, consequently, as incomplete as the DM's understanding of the problem.

In the Holmes-Watson example, redefine a as the state in which there is a human intruder, there is no canine intruder, and Watson knows it; b the state in which there is a canine intruder, there is no human intruder, and Watson knows it; and c the state in which there is no intruder at all and Watson is unaware of both the possibility of there being a canine or a human intruder. It turns out that we need three further states to describe Watson's unawareness: the state at which Watson is unaware of both possibilities, there is no canine intruder and there is a human intruder (d); the state at which Watson is unaware of both possibilities, there is a canine intruder and no human intruder (e); and finally, the state at which Watson is unaware of both possibilities (f). These states are necessary to describe Watson's epistemic status for reasons that should become clear shortly. Table 7 summarizes the list of states; in it, H and C are taken to be the propositions "there was a human intruder" and "there was a canine intruder," respectively.

ω	H	$\neg H$	C	$\neg C$	$k(H)$	$k(\neg H)$	$k(C)$	$k(\neg C)$	$u(H)$	$u(C)$
a	×			×	×			×		
b		×	×			×	×			
c		×		×					×	×
d	×			×					×	×
e		×	×						×	×
f									×	×

Table 7: The Expanded State Space

Observe that while the description of states a,b,c,d and e entails the description of both the possibility of a canine and of a human intruder, the description of state f does not involve any possibility of a canine or a human intruder. In that sense, these are really two different groups of states, as they require different expressive powers. Hence, the state space can be regarded as comprising the states from the spaces $\Omega_{C,H} = \{a,b,c,d,e\}$ (where the possibilities of a canine and a human intruder can be expressed), and $\Omega_{\varnothing} = \{f\}$ (where no possibility can be expressed). Intuitively, the DM "lives" in one such space.

The event "there was a canine intruder" corresponds to $\{b,e\}$; the event "there was not a canine intruder" to $\{a,c,d\}$. Observe that the state f belongs to neither event, in spite of the fact that one is the complement of the other. This is so because this state belongs to a space where the possibility of a canine intruder cannot be described. As remarked earlier, the state is no longer a complete description of all relevant variables, as in Savage, but a complete description of all relevant variables according to the DM, at his present epistemic status, which may be incomplete.

As a consequence, in general we now have $E \cup E^c \quad \Omega$ for those events related to the DM's environment. Moreover, complements should be considered in the space where the event is defined (and upwards), for the proposition that defines it cannot be described in lower-expressive power spaces. Schipper, then, defines events as the states in the base space that satisfy a proposition <u>plus</u> their projection into higher-expressive power spaces.

Interestingly, the universe and the empty sets from all these sets should be regarded as different. Recall that every set is the collection of elements that satisfy a given proposition, and hence the empty set represents a contradiction and the universe a tautology. The point to note is that in all the spaces different propositions can be expressed, and hence different tautologies and contradictions. We will thus label $\varnothing_{C,H}, \varnothing_{\varnothing}$ the contradictions in each of them. Being aware of $\varnothing_{C,H}$ is logically distinct from $\varnothing_{\varnothing}$, for the former implies awareness of the possibility of a canine and a human intruder, while the latter of none.

In contrast, all states specify the DM's epistemic condition, so that for those events regarding his epistemic status we have $E \cup E^c = \Omega$, as usual.

In the example, we have $\Omega = \{a,b,c,d,e,f\}$. Watson's possibility correspondence may be represented as in Table 7.

ω	$P_W(\omega)$
a	$\{a\}$
b	$\{b\}$
c	$\{c,d,e,f\}$
d	$\{c,d,e,f\}$
e	$\{c,d,e,f\}$
f	$\{c,d,e,f\}$

Table 8: Watson's Possibility Correspondence in Schipper's View

From this possibility correspondence we can readily obtain Watson's knowledge structure. Table 9 reproduces it for a selected number of events.

E	$K(E)$	$\neg K(E)$	$K(K(E))$	$K(\neg K(E))$	$\neg K(\neg K(E))$	$U(E)$
$\{a\}$	$\{a\}$	$\{b,c,d,e\}$	$\{a\}$	$\{b\}$	$\{a,c,d,e\}$	$\{c,d,e\}$
$\{b\}$	$\{b\}$	$\{a,c,d,e\}$	$\{b\}$	$\{a\}$	$\{b,c,d,e\}$	$\{c,d,e\}$
$\{c\}$	$\varnothing_{C,H}$	$\{a,b,c,d,e\}$	$\varnothing_{C,H}$	$\{a,b\}$	$\{c,d,e\}$	$\{c,d,e\}$
$\{c,d,e\}$	$\varnothing_{C,H}$	$\{a,b,c,d,e\}$	$\varnothing_{C,H}$	$\{a,b\}$	$\{c,d,e\}$	$\{c,d,e\}$
$\{c,d,e,f\}$	$\{c,d,e,f\}$	$\varnothing_{\varnothing}$	$\{c,d,e,f\}$	$\varnothing_{\varnothing}$	$\{c,d,e,f\}$	$\varnothing_{\varnothing}$
$\{b,e\}$	$\{b\}$	$\{a,c,d,e\}$	$\{b\}$	$\{a\}$	$\{b,c,d,e\}$	$\{c,d,e\}$

Table 9: Watson's Knowledge Structure in Schipper's View

Two elements are noteworthy. Firstly, observe that $U\left(U\left(\{c,d,e\}\right)\right)=U\left(\{c,d,e\}\right)$, that is, in the states in which the DM is unaware of an event, he is also unaware of this fact. Since this is a general property of the model, it is immune to the DKL critique. Secondly, compare the event $\{c,d,e\}$, "Watson is unaware of the dog and the man" with the event $\{c,d,e,f\}$, "Watson is unaware of something." Watson is unaware of the former, but never of the latter.

Overall, Schipper's model is certainly promising as a "grammar" for unawareness that, as the reader may have noticed, is based on sets rather than propositions. Li (2003) independently developed a model along the same basic lines, namely, incorporating in the state the description of the DM's epistemic status, and allowing for "personal" (thus incomplete) states. Li's model is also based on sets. This is important because set-based models are much more tractable, as the literature following Aumann (1976) could appreciate.

On top of this grammar, however, we need a model of behavior before we can analyze behavior under unawareness. We currently lack such a model, but good progress has been made, and it seems very close as Section 4 relates.

3.2 The Behavioral Consequences of Unawareness

Information transmission in financial markets has been studied theoretically ever since the development of information economics. While some authors sought to describe in a model a situation where investors would trade based on disparate information (Lintner 1969, Hirshleifer 1977) generating divergent opinions over capital gains, or speculation, economists promptly abandoned this idea on the grounds that rationality would preclude it as an equilibrium phenomenon. Thus, Radner (1979) showed that if investors knew the connection between asset prices and states, then generic equilibrium prices would be perfectly revealing. Milgrom and Stokey (1982) extended the result to more general trading environments. Subsequent research has shown this conclusion to be fairly robust even in the absence of common knowledge of rationality and other troublesome assumptions.

Heifetz, Meier and Schipper (2003) show an example, however, where two traders agree to trade based on private information. This can occur because each of them is unaware of the fact that her trading partner also possesses information, and furthermore, each holds a theory of her opponent's behavior which is consistent with what she observes.

Hence, the development of a theory of unawareness revives the old question of information transmission in markets. Since Hayek (1945), economists have sought to understand modern economies' division of labor in the task of information processing, yet the rational expectations revolution somehow froze the question.

On a different subject, it seems natural that two parties entering into a contract may be unaware of some fact that, when revealed, leads them to revise their previous decisions.

The balance between commitment and *ex post* optimality is a delicate one in this context. While the purpose of a contract is to commit to particular, *ex post* undesirable actions, this commitment is usually made on limited information. The current literature assumes a commonly known state space, precluding unawareness. Since the contract may condition on the state, the initial beliefs (probability) do not affect the optimal contract. Explicitly considering unawareness forces us to revise these conclusions.

On the other hand, a new benefit of the contract arises under this fresh look: the contract itself may generate awareness of certain possibilities –and hence there could be strategic incompleteness: a party aware of an event but believing his counterpart unaware, does not bring the subject into the discussion to maintain his current state of mind (e.g., the person offering a house for rent does not make the contract conditional on plans to build a nuclear plant nearby).

These are just a few (important) examples from the myriad of social phenomena that we would be able to revisit with an expanded Savage paradigm.

4. Non-expected Utility

4.1 Modeling ambiguity

Ellsberg (1961) suggested two thought experiments that would be critical to the predictive value of EUT under any interpretation. He conjectured ordinary people would behave in a manner that would be negative for the theory, and subsequent experimental work proved him right (Camerer, 1995, pg. 646, surveys some of this experimental work).

There are many versions of the paradox, but all of them are variations on the following idea. Consider an individual who is presented two urns (A and B) with 100 chips each, either red (r) or blue (b). The individual is told that urn A contains exactly 50 red and 50 blue chips, and that urn B also contains only red and blue chips, but he is not told in which proportions. The individual is asked to bet on a color: one chip is extracted, and if it matches the color of his choosing, he wins a prize of value x. If the color is different, he gets 0. Before betting on the color, however, he is asked to choose the urn from which the chip will be extracted.

Most individuals in experiments of this sort say they do not care which color they bet on, but prefer to bet on urn A, the one with the known composition. This behavior is inconsistent with the EU model. To see this, let us compute the expected utilities of each alternative. The following list indicates the expected utility of betting on urns A and B and colors r and b:

$$
\begin{aligned}
E[u](A;r) &= \pi(r_A)u(x) + \pi(b_A)u(0) \\
E[u](A;b) &= \pi(r_A)u(0) + \pi(b_A)u(x) \\
E[u](B;r) &= \pi(r_B)u(x) + \pi(b_B)u(0) \\
E[u](B;b) &= \pi(r_B)u(0) + \pi(b_B)u(x)
\end{aligned} \tag{14}
$$

If the individual is indifferent to the colors in both urns, he must associate a 50% chance (whatever that means) to obtaining each color from either urn:

$$
\begin{aligned}
E[u](A;r) = E[u](A;b) &\Rightarrow \pi(r_A) = \pi(b_A) = \frac{1}{2} \\
E[u](B;r) = E[u](B;b) &\Rightarrow \pi(r_B) = \pi(b_B) = \frac{1}{2}
\end{aligned} \tag{15}
$$

This implies, however, that he must associate the same utility level to both urns:

$$
E[u](A) = \frac{1}{2}\left[u(x) + u(0)\right] = E[u](B). \tag{16}
$$

Taken from a different perspective, indifference to colors in A means:

$$
E[u](A) = \frac{1}{2}\left[u(x) + u(0)\right] \tag{17}
$$

Meanwhile, the utility of choosing urn B is given by:

$$
E[u](B) = \max\left\{\pi(r_B)u(x) + \pi(b_B)u(0); \pi(r_B)u(0) + \pi(b_B)u(x)\right\} \tag{18}
$$

so that if $\pi(r_B) < \frac{1}{2}$, the individual would prefer to bet on blue and

$$
E[u](B) = \pi(r_B)u(0) + \pi(b_B)u(x) > \frac{1}{2}\left[u(x) + u(0)\right] = E[u](A) \tag{19}
$$

If $\pi(r_B) > \frac{1}{2}$, the individual would rather bet on red but still on urn B, using the same argument. Therefore, there is no scenario we can think of in which urn A is preferred to urn B if the individual associates a 50% chance with each color in A. If he is indifferent to colors in B, he must be indifferent to which urn is used. If he is not indifferent to the colors in B, then he must prefer B.

Hence, if an EU-maximizer is indifferent to the colors in both urns, he must be indifferent to which urn is used, a prediction systematically contradicted by experimental data. It is important to stress that this prediction is fully behavioral, and does not depend on the analyst's or DM's understanding of what a probability is, or even on whether a EU-maximizer thinks consciously or not.

This section describes one effort to amend the EU model to ensure it describes behavior of this nature, namely, the Choquet Expected Utility model (henceforth, CEU.) The chief proposal is to consider a generalization of probabilities – called *capacities* – to represent beliefs, thus giving rise to an expected utility representation of preferences where the probabilities are not additive. There are other theories that rationalize non-additive beliefs as well. CEU is possibly the most well known, but this is an area of active research, and new developments are to be expected in the near future.

As with the EU model, many interpretations can be given to the CEU model. Originally, adherents to objective interpretations of probability coined the term "Knightian uncertainty" to refer to situations in which the decision-maker does not *know* the probabilities of events. The term derives from a distinction between risk (the DM cannot tell in advance what the outcome of an experiment will be, but *knows* the probabilities of each outcome, as in urn A) and uncertainty (where the *knowledge* of these probabilities is absent, as in urn B) attributed to Frank Knight (1921). For those authors, the CEU model is adequate to describe Knightian uncertainty.

This distinction does not make sense under a subjective interpretation, for probabilities represent the beliefs of the DM built into his decision criterion together with his tastes, and hence whether the individual is conscious of his own doubts or not is of no consequence (recall the anecdote of the classroom discussion.) All that matters is the (implicit) degree of trust in the occurrence of each event, which rationalizes his choices.

A subjectivist, on the other hand, can also make sense of this model. Under a subjective interpretation of probability, Savage's axioms constitute a nice argument in favor of having EU-representable behavior, but not a proof – and the proof must lie in actual behavior. Hence, any systematic violation of EU-represented behavior is evidence against some of Savage's axioms. A subjectivist, then, may momentarily assume that beliefs have a representation as probabilities, but cannot prove it, for the proof is in the observed behavior. Ellsberg's paradox is an example of a systematic departure from EU-representable behavior, and Schmeidler's (1989) axiomatization, which results in a CEU representation, offers a plausible adaptation of Savage's work to accommodate it.

Perhaps with the purpose of welcoming subjectivists to the study of non-additive beliefs, or maybe because some subjectivists took it over, lately the early term "Knightian uncertainty" has been replaced by "ambiguity," which seems inoffensive to either view.

Explaining comparatively the axiomatizations of each theory goes beyond the purpose of this essay. Instead, this section attempts to explain at an intuitive level how CEU works. The reader who is interested in the formal structure of the model will be best served by going to the sources.

In the context of the example, Schmeidler's observation is that we may disassociate indifference to colors – no reason to prefer one color over another – from indifference to urns. Instead of associating a "probability" to each event, let us say that the individual associates a degree of confidence to the occurrence of a state, v, not necessarily represented by a probability. In particular, rewrite Equation 14 as:

$$\begin{aligned}
E[u](A;r) &= v(r_A)u(x) + (1 - v(r_A))u(0) \\
E[u](A;b) &= (1 - v(b_A))u(0) + v(b_A)u(x) \\
E[u](B;r) &= v(r_B)u(x) + (1 - v(r_B))u(0) \\
E[u](B;b) &= (1 - v(b_B))u(0) + v(b_B)u(x)
\end{aligned} \tag{20}$$

Indifference to colors implies:

$$\begin{aligned}
E[u](A;r) = E[u](A;b) \Rightarrow v(r_A) = v(b_A) \\
E[u](B;r) = E[u](B;b) \Rightarrow v(r_B) = v(b_B)
\end{aligned} \tag{21}$$

However, urn A is preferred to urn B as long as $v(r_A) > v(r_B)$:

$$\begin{aligned}
&E[u](A) = v(r_A)u(x) + (1 - v(r_A))u(0) \\
&> E[u](B) = v(r_B)u(x) + (1 - v(r_B))u(0)
\end{aligned} \tag{22}$$

The interpretation given to the $v(\bullet)$ function is that it represents both a degree of confidence in the occurrence of an event and a measure of the ambiguity that the DM perceives in his decision problem. Urn A represents a less ambiguous choice than urn B, because the individual has more information and hence more confidence in his beliefs, even though he has not reason to believe that one color is more likely than the other for either urn.

The difference with respect to expected utility theory is that $v(r_A) + v(r_B) \neq 1$, that is, the belief is not additive. Mathematically, the belief is not represented by a probability function but by a capacity. A capacity is a non-additive probability, that is, a function $v : \mathfrak{M} \to [0,1]$ that satisfies $\pi 1 - \pi 2$ but not $\pi 3$ (in Section 2.1). Thus, a probability is an additive capacity.

Schmeidler (1989) proves that an EU-representation of preferences still exists when one of Savage's axioms is relaxed, but where the expectation is not taken over a probability but over a capacity. The lack of additivity, however, implies that the usual integral cannot be applied here. The appropriate integral concept is that of Choquet (which explains the name of the model). Let $f(C)=\{c_1,c_2,...,c_n\}$ be the set of consequences under act f, where $u(c_i)\geq u(c_{i+1})$ for all i. The Choquet integral is given by:

$$\int u(f)dv=\sum_{i=1}^{n-1}\left[u\left(c_i-u\left(c_{i+1}\right)\right)\right]v\left(\bigcup_{j=1}^{i}A_j\right)+u(c_n) \tag{23}$$

where $A_j=f^{-1}(c_j)$ is the event in which consequence c_j obtains under act f.

In the two-state urn example that introduced this section, this definition means that as long as $v(r_\bullet)+v(b_\bullet)<1$, the higher-utility scenario is weighted by $v(\bullet)$, the lower-utility scenario by $1-v(\bullet)$, that is, the degree of confidence in its occurrence **plus** all non-assigned weight.

If there were two events, this would imply that the indifference curves over risky consumption profiles are kinked (and hence non differentiable) at the certainty line. For instance, let $\Omega=\{\omega_1,\omega_2\}$ be the state space and the set of consequences $\mathbb{R}^2$, that is, ordered pairs of positive reals. Then, each act is a bundle (c_1,c_2) indicating a consumption level for each state. Suppose beliefs have the form given by Table 10.

E	$v(E)$
$\varnothing$	0
$\{\omega_1\}$	$\frac{1}{4}$
$\{\omega_2\}$	$\frac{1}{4}$
$\{\omega_1,\omega_2\}$	1

Table 10: A Convex Capacity

Then, if f is such that $c_1>c_2$,

$$\int u(f)dv=\frac{1}{4}u(c_1)+\frac{3}{4}u(c_2)$$

whereas in the opposite case, when $c_1<c_2$

$$\int u(f)dv=\frac{3}{4}u(c_1)+\frac{1}{4}u(c_2)$$

The Choquet integral, then, adds the unassigned weight $1-v(E)-v(E^c)$ to the worst possible outcome. The certainty line separates the cases where state ω_1 is associated to the worst outcome from the cases where it is state ω_2.

A capacity is said to be convex if it satisfies:

$$v(A\cup B)+v(A\cap B)\geq v(A)+v(B) \quad \text{for all } A,B\in\mathfrak{M} \tag{24}$$

Hence, for a convex capacity, $1\geq v(A)+v(A^c)$ for all $A\in\mathfrak{M}$, as in the urn example. However, this need not be always the case. Consider the following example:

E	$v(E)$
$\varnothing$	0
$\{\omega_1\}$	$\frac{3}{4}$
$\{\omega_2\}$	$\frac{3}{4}$
$\{\omega_1,\omega_2\}$	1

Table 11: A Concave Capacity

The CEU in this case is given by:

$$\int u(f)dv=\begin{cases}\dfrac{3}{4}u(c_1)+\dfrac{1}{4}u(c_2) & \text{if } c_1>c_2\\[2mm] \dfrac{1}{4}u(c_1)+\dfrac{3}{4}u(c_2) & \text{otherwise}\end{cases} \tag{25}$$

Observe that a DM with such beliefs assigns the "extra" weight (the excess over 1) to the best possible outcome, that is, precisely the opposite of the convex-capacity DM. Intuitively, the former is being "optimistic", while the latter "pessimistic." Recall that the form of the capacity is obtained by looking at the DM's behavior, so beliefs are not really separable from attitudes towards ambiguity, as the attitudes towards risk case. For that reason, a DM with a convex capacity is called "ambiguity averse" (and the degree of his aversion measured by the difference $1-\left[v(A)+v(A^c)\right]$), while a DM with a concave capacity is called an "ambiguity lover", paralleling the terminology "risk averse-risk lover."

A very interesting feature of this model is that it is capable of rationalizing puzzling behavior beyond Ellsberg's thought experiments. Consider for instance the case of a person who bets heavily at a casino and at the same time buys life and health insurance. While a risk-averse EU-maximizer is risk-averse in every dimension, a CEU-maximizer need not hold equally non-additive beliefs regarding every uncertainty. Thus, the same individual can have non-additive beliefs with respect to one variable and additive with respect to another.

Imagine a person who entertains the following possibilities for the next day: {it rains-there is an earthquake, it rains-there isn't an earthquake, it doesn't rain-there is an earthquake, it doesn't rain-there isn't an earthquake} $\equiv \{RE, RN, DE, DN\}$. The capacity of Table 12 is additive with regard to the possibility of rain, but non-additive with regard to the possibility of an earthquake (the 3-state events are omitted for brevity):

E	$v(E)$	E	$v(E)$
$\varnothing$	0	$\{RE, RN\}$	0.7
$\{RE\}$	0.07	$\{DE, DN\}$	0.3
$\{RN\}$	0.1	$\{RE, DE\}$	0.2
$\{DE\}$	0.03	$\{RN, DN\}$	0.3
$\{DN\}$	0.1	$\{RE, RN, DE, DN\}$	1

Table 12: A Capacity that is Non-Additive in One Dimension while Additive in the Other

In this example, the individual associates a probability of 70% to rain and 30% to not rain, and therefore he will behave as an EU maximizer if faced with decisions whose consequences depend exclusively on whether it will rain or not. However, his degree of confidence in an earthquake happening is 20%, it not happening 30%, and hence he will behave as a CEU maximizer in decisions involving earthquake/no earthquake as a state. The remaining 50% can be interpreted as the degree to which the situation seems ambiguous to him. While this person will behave as very risk averse in situations that depend on the occurrence of an earthquake, he will not do so in situations that depend on the rain. Hence, the generalization presented by CEU is non-trivial.

The areas in which CEU theory yields interesting predictions are abundant. One example is offered by the non-participation puzzle. While EUT predicts that perfectively competitive investors would hold every available asset in their portfolios (though in different proportions, accounting for differences in beliefs, risk aversion or endowments), typical households do not hold stocks. For instance, in the US only one of every five households does. In contrast, Dow and Werlang (1992) showed that if an investor had (convex) non-additive beliefs regarding assets' returns, then there would be a range of prices at which he would not buy or sell, but stay out of the market. Ambiguity aversion offers a rationalization of this puzzle.

In contrast to unawareness, ambiguity is essentially a behavioral phenomenon. Consequently, there is a much richer literature exploring its implications in social situations. Marinacci (2000), for instance, extends the theory of normal form games to players with ambiguous beliefs. Eichberger, Kelsey and Schipper (2003) revisit classic problems, like that of oligopoly in the views of Cournot and Bertrand. In most cases bringing ambiguity into consideration significantly alters individual behavior and equilibrium values as well.

5. Further Remarks

The development of the EU-model and its associated theories is possibly one of the greatest intellectual achievements of the 20th century. During the second half, it made possible the development of entire fields, among them financial economics, game theory, the economics of information, and modern macroeconomics. Some questions still remain, however, that it has not proved capable of answering in its present form, particularly in situations in which the uncertainty goes "beyond earthquakes and the weather."

This essay discussed separately two aspects of individual decision making with potentially important behavioral content: unawareness and ambiguity. It turns out that these may be intimately connected, as Ghirardato's (2001) research suggests.

Ghirardato adapted Savage's model to the unawareness case. His construction consists of separating the individual's subjective state space from the objective state space. The former reflects the DM's understanding of the problem, while the latter the true possibilities. Given that the DM's understanding is limited, his state space is coarser than the objective space. Thus, the consequence of an act $c(a,\omega)$ is unknown even when conditional on a state; mathematically, $c(a,\omega)$ it is a set instead of a point. After extending Savage's axioms he obtains the same EU representation of preferences and therefore a probabilistic representation of beliefs, but over the subjective space. Adding some more structure he manages to obtain another representation, this time over the objective state space, which turns out to be non-additive. Hence, an observer might be able to represent a subject's behavior by means of a CEU *because of* the subject's unawareness.

At any rate, these extensions of EUT are an active research field, which is likely to produce new answers to old problems. Similarly, the associated development of new tools will make economic theory suitable for the analysis of new problems as well, expanding its frontiers, just as EUT itself did in the past.

References

Allen, F. and D. Gale (1989), "Optimal Security Design," The Review of Financial Studies **1**, 229-263.

Anscombe, F. and R. Aumann (1963), "A Definition of Subjective Probability," Annals of Mathematical Statistics **34**, 199-205.

Aumann, R. (1976), "Agreeing to Disagree," Annals of Statistics **4**, 1236-1239.

Aumann, R. (1998), "Common Priors: A Reply to Gul," Econometrica **66**, 929-938.

Aumann, R. (1999a), "Interactive Epistemology I: Knowledge," International Journal of Game Theory **28**, 263-300.

Aumann, R. (1999a), "Interactive Epistemology II: Probability," International Journal of Game Theory **28**, 301-314.

Arrow, K., and F. Hahn (1999), "Notes on Sequence Economies, Transaction Costs and Uncertainty," Journal of Economic Theory **86**, 203-218.

Bernardo, J. and A. Smith (1994), "Bayesian Theory," Wiley.

Bernoulli, D. (1738), "Exposition of a New Theory on the Measurement of Risk," reprinted and translated in Econometrica **22**, 23-36, 1954.

Camerer, C. (1995), "Individual Decision Making," in The Handbook of Experimental Economics edited by J. Hagel and A. Roth, Princeton.

Dekel, E., B. Lipman, and A. Rustichini (1998a), "Standard State-Space Models Preclude Unawareness," Econometrica **66**, 159-173.

Dekel, E., B. Lipman, and A. Rustichini (1998b), "Recent Developments in Modeling Unforeseen Contingencies," European Economic Review **42**, 523-542.

Dow, J. and Werlang, S.R. (1992), "Uncertainty Aversion, Risk Aversion, and the Optimal Choice of Portfolio," Econometrica **60**, 197-204.

Eichberger, J., D. Kelsey and B. Schipper (2003), "Ambiguity and Social Interaction," mimeo.

Ellsberg, D. (1961), "Risk, Ambiguity, and the Savage Axioms," Quarterly Journal of Economics **75**, 643-669.

Ely, J. (1998), "A Note on Unawareness," mimeo.

Friedman, M. (1953), “The Methodology of Positive Economics,” in Essays in Positive Economics, Chicago.

Geanakoplos, J. (1989), “Game Theory without Partitions, and Applications to Speculation and Consensus,” Cowles Foundation Discussion Paper No. 914.

Ghirardato, P. (1993), “Non-Additive Measures of Uncertainty: A Survey of Some Recent Developments in Decision Theory,” International Review of Economics and Business **40**, 253-276.

Ghirardato, P. (2001), “Coping with Ignorance: Unforeseen Contingencies and Non-additive Uncertainty,” Economic Theory **17**, 247-276.

Grossman, S. and J. Stiglitz (1980), “On the Impossibility of Informationally Efficient Markets,” American Economic Review **70**, 393-408.

Gul, F. (1998), “A Comment on Aumann’s Bayesian View,” Econometrica **66**, 923-927.

Hart, O. and J. Moore (1988), “Incomplete Contracts and Renegotiation,” Econometrica **56**, 755-785.

Hájek, A. (2002), “Interpretations of Probability,” Stanford Encyclopedia of Philosophy (available at http://plato.stanford.edu)

Hayek, F. (1945), “The use of knowledge in society,” American Economic Review **35**, 519-530.

Heifetz, A., M. Meier and B. Schipper (2003), “Interactive Unawareness and Speculative Trade,” mimeo.

Hirshleifer, J. (1977), “The Theory of Speculation under Alternative Regimes of Markets,” Journal of Finance **32**, 975-999.

Knight, F. (1921), “Risk, Uncertainty and Profit.”

Kolmogorov, A. (1950), “Foundations of the Theory of Probability,” Chelsea.

Kreps, D. (1988), “Notes on the Theory of Choice,” Westview.

Li (2003), “Unawareness,” mimeo.

Lintner, J. (1969), “The Aggregation of Investors’ Diverse Judgments and Preferences in Purely Competitive Security Markets,” Journal of Financial and Quantitative Analysis **4**, 103-124.

Marinacci, M. (2000), “Ambiguous Games,” Games and Economic Behavior **31**, 191-219.

Maskin, E. and J. Tirole (1999), “Unforeseen Contingencies and Incomplete Contracts,” Review of Economic Studies **66**, 83-114.

McGoun, E. (1995), “The History of Risk Measurement,” Critical Perspectives on Accounting **6**, 511-532.

Milgrom, P. and N. Stokey (1982), “Information, Trade, and Common Knowledge,” Journal of Economic Theory **26**, 17-27.

Morris, S. (1995), “The Common Prior Assumption in Economic Theory,” Economics and Philosophy **11,** 227-253.

Osborne, M. and A. Rubinstein (1994), “A Course in Game Theory,” MIT.

Radner, R. (1979), “Rational Expectations Equilibrium: Generic Existence and the Information Revealed by Prices,” Econometrica **47**, 655-678.

Riley, J. (1979), “Informational Equilibria,” Econometrica **47**, 331-360.

Savage, L. (1957), “The Foundations of Statistics,” Wiley. Revised edition, Dover, 1972.

Schipper, B. (2002), “A State Space Structure with Unawareness,” mimeo.

Spence, M. (1973), “Job Market Signalling,” Quarterly Journal of Economics **87**, 296-332.

Schmeidler, D. (1989), “Subjective Probability and Expected Utility without Additivity,” Econometrica **57**, 571-587.

Vickrey, W. (1961), “Counterspeculation, Auctions, and Competitive Sealed Tenders,” Journal of Finance **16**, 8-37

Von Neumann, J. and O. Morgenstern (1946), “The Theory of Games,” Princeton.