Game Theory without Decision-Theoretic Paradoxes

Pierfrancesco La Mura
Department of Economics and Information Systems
HHL - Leipzig Graduate School of Management
Jahnallee 59, 04109 Leipzig (Germany)
plamura@hhl.de

Abstract

Most work in game theory is conducted under the assumption that the players are expected utility maximizers. Expected utility is a very tractable decision model, but is prone to well-known paradoxes and empirical violations (Allais 1953, Ellsberg 1961), which may induce biases in game-theoretic predictions. La Mura (2009) introduced a projective generalization of expected utility (PEU) which avoids the dominant paradoxes, while remaining quite tractable. We show that every finite game with PEU players has an equilibrium, and discuss several examples of PEU games.

Keywords: game theory, equilibrium, expected utility, paradoxes, Allais, Ellsberg, projective

1 Introduction

The expected utility hypothesis is the *de facto* foundation of game theory.

The von Neumann - Morgenstern axiomatization of expected utility, and later on the subjective formulations by Savage (1954) and Anscombe and Aumann (1963) were immediately greeted as simple and intuitively compelling. Yet, in the course of time, a number of empirical violations and paradoxes (Allais 1953, Ellsberg 1961) came to cast doubt on the validity of the hypothesis as a foundation for the theory of rational decisions in conditions of risk and subjective uncertainty. In economics and in the social sciences, the shortcomings of the expected utility hypothesis are generally well-known, but often tacitly accepted in view of the great tractability and usefulness of the corresponding mathematical framework. In fact, the hypothesis postulates that preferences can be represented by way of a utility functional which is linear in probabilities, and linearity makes expected utility representations particularly tractable in models and applications. In particular, in a game-theoretic context, linearity of expected utility ensures that the best response correspondence is also linear, and hence that any finite game has a Nash equilibrium.

In economics and the social sciences, the importance of accounting for violations of the expected utility hypothesis has long been recognized (Tversky 1975), but so far none of its numerous alternatives (e.g., Machina 1982, Schmeidler 1989, to quote only two particularly influential papers in a rich and constantly evolving literature) has fully succeeded in replacing expected utility as a standard foundation for decisions under uncertainty, partly due to the great mathematical tractability of expected utility relative to many of its proposed generalizations.

La Mura (2009) introduced projective expected utility (PEU), a decision-theoretic framework which accommodates the dominant paradoxes while retaining significant simplicity and tractability. This is obtained by weakening the expected utility hypothesis to its projective counterpart, in analogy with the quantum-mechanical generalization of classical probability theory.

We first review the EU and PEU frameworks, and show that the latter is sufficiently general to avoid both Allais' and Ellsberg's paradoxes. We then extend the notion of Nash equilibrium to games with PEU preferences, and prove that any finite game with PEU players has an equilibrium. Finally, we discuss several examples of games with PEU preferences and identify observable deviations from the theory of games with EU preferences.

2 von Neumann - Morgenstern Expected Utility

Let Ω be a finite set of outcomes, and Δ be the set of probability functions defined on Ω, taken to represent risky prospects (or *lotteries*). Next, let $\succsim$ be a complete and transitive binary relation defined on $\Delta \times \Delta$, representing a decision-maker's preference ordering over lotteries. Indifference of $p, q \in \Delta$ is defined as $[p \succsim q$ and $q \succsim p]$ and denoted as $p \sim q$, while strict preference of p over q is defined as [$p \succsim q$ and not $q \succsim p$], and denoted by $p \succ q$. The preference ordering is assumed to satisfy the following two conditions.

Axiom 1 *(Archimedean) For all $p, q, r \in \Delta$ with $p \succ q \succ r$, there exist $\alpha, \beta \in (0, 1)$ such that $\alpha p + (1 - \alpha)r \succ q \succ \beta p + (1 - \beta)r$.*

Axiom 2 *(Independence) For all $p, q, r \in \Delta$, $p \succsim q$ if, and only if, $\alpha p + (1 - \alpha)r \succsim \alpha q + (1 - \alpha)r$ for all $\alpha \in [0, 1]$.*

A functional $u : \Delta \to \mathbb{R}$ is said to represent $\succsim$ if, for all $p, q \in \Delta$, $p \succsim q$ if and only if $u(p) \geq u(q)$.

Theorem 1 *(von Neumann and Morgenstern) Axioms 1 and 2 are jointly equivalent to the existence of a functional $u : \Delta \to \mathbb{R}$ which represents $\succsim$ such that, for all $p \in \Delta$,*

$$u(p) = \sum_{\omega \in \Omega} u(\omega)p(\omega).$$

The von Neumann - Morgenstern setting is appropriate whenever the nature of the uncertainty is purely objective: all lotteries are associated with objective random devices, such as dice or roulette wheels, with well-defined and known frequencies for all outcomes, and the decision-maker only evaluates a lottery based on the frequencies of its outcomes.

3 Allais' Paradox

The following paradox is due to Allais (1953). First, please choose between

A: A chance of winning 4000 dollars with probability 0.2

B: A chance of winning 3000 dollars with probability 0.25.

Now suppose that, instead of A and B, your two alternatives are

C: A chance of winning 4000 dollars with probability 0.8

D: A chance of winning 3000 dollars with certainty.

If you chose A over B, and D over C, then you are in the modal class of respondents. The paradox lies in the observation that A and C are special cases of a two-stage lottery E which in the first stage either returns zero dollars with probably $(1 - \alpha)$ or, with probability α, leads to a second stage where one gets 4000 dollars with probability 0.8 and zero otherwise. In particular, if α is set to 1 then E reduces to C, and if α is set to 0.25 it reduces to A. Similarly, B and D are special cases of a two-stage lottery F which again with probability $(1 - \alpha)$ returns zero, and with probability α continues to a second stage where one wins 3000 dollars with probability 1. Again, if $\alpha = 1$ then F reduces to D, and if $\alpha = 0.25$ it reduces to B. Then it is easy to see that the [$A \succ B$, $D \succ C$] pattern violates Axiom 2 (Independence), as E can be regarded as a lottery $\alpha p+(1-\alpha)r$, and F as a lottery $\alpha q+(1-\alpha)r$, where p and q represent lottery C and D, respectively, and r represents the lottery in which one gets zero dollars with certainty. When comparing E and F, why should it matter what is the value of α? Yet, experimentally one finds that it does.

4 Ellsberg's Paradox

Another disturbing violation of the expected utility hypothesis was pointed out by Ellsberg (1961). Suppose that an urn contains 300 balls of three possible colors: red, green, and blue. You know that the urn contains exactly 100 red balls, but are given no information on the proportion of green and blue.

You win if you guess which color will be drawn. Do you prefer to bet on red (R) or on green (G)? Many respondents choose R, on grounds that the probability of drawing a red ball is known to be $1/3$, while the only information on the probability of drawing a green ball is that it is between 0 and $2/3$. Now suppose that you win if you guess which color will *not* be drawn. Do you prefer to bet that red will not be drawn ($\overline{R}$) or that green will not be drawn ($\overline{G}$)? Again many respondents prefer to bet on $\overline{R}$, as the probability is known ($2/3$) while the probability of $\overline{G}$ is only known to be between 1 and $1/3$.

The pattern [$R \succ G$, $\overline{R} \succ \overline{G}$] is incompatible with von Neumann - Morgenstern expected utility, which only deals with known probabilities, and is also incompatible with the Savage (1954) formulation of expected utility with subjective probability as it violates its Sure Thing axiom. Observe that, in Ellsberg's setting, the decision-maker ignores the actual composition of the urn, and hence operates in a state of subjective uncertainty about the true probabilistic state of affairs. In particular, the decision-maker is exposed to a combination of subjective uncertainty (on the actual composition of the urn) and objective risk (the probability of drawing a specific color from an urn of given composition). The paradox suggests that, in order to account for the choice pattern discussed above, subjective uncertainty and risk should be handled as distinct notions.

5 Projective Expected Utility

Let X be the positive orthant of the unit sphere in $\mathbb{R}^n$, where n is the cardinality of the set of relevant outcomes Ω. Then von Neumann - Morgenstern lotteries, regarded as elements of the unit simplex, are in one-to-one correspondence with elements of X, which can therefore be interpreted as risky prospects, for which the frequencies of the relevant outcomes are fully known. Observe that, while the projections of elements of the unit simplex (and hence, L^1 unit vectors) on the basis vectors can be naturally associated with probabilities, if we choose to model von Neumann - Morgenstern lotteries as elements of the unit sphere (and hence, as unit vectors in L^2) then probabilities are naturally associated with *squared* projections. The advantage of such move is that L^2 is the only L^p space which is also a Hilbert space, and Hilbert spaces have a very tractable projective structure which is exploited by the representation. In particular, it is unique to L^2 that the set of unit vectors is invariant with respect to projections.

Next, let $\langle .|. \rangle$ denote the usual inner product in $\mathbb{R}^n$. We denote the transpose of a vector or a matrix with a prime, *e.g.*, x' denotes the transpose of x. An orthonormal basis is a set of unit vectors $\{b^1, ..., b^n\}$ such that $\langle b^i | b^j \rangle = 0$ whenever $i \neq j$. In our context, orthogonality captures the idea that two events or outcomes are mutually exclusive (for one event to have probability one, the other must have probability zero). The natural basis corresponds to the set of degenerate lotteries returning each objective lottery outcome with certainty, and is conveniently identified with the set of objective lottery outcomes $\{\omega^1, \omega^2, ..., \omega^n\}$. Yet, in any realistic experimental setting, it is very unlikely that those objective outcomes will happen to coincide with the set of subjective consequences which are relevant from the perspective of the decision-maker. Moreover, even if the latter could be fully elicited, it would be generally problematic to relate a von Neumann - Morgenstern lottery, which only specifies the probabilities of the objective outcomes, with the probabilities induced on the subjective consequences, that is, the relevant dimensions of risk from the point of view of the decision-maker. The perspective of the observer or modeler is inexorably bound to objectively measurable entities, such as frequencies and prizes; by contrast, the decision-maker thinks and acts based on subjective preferences and subjective consequences, which in a revealed-preference context should be presumed to exist while at the same time assumed, as a methodological principle, to be unaccessible to direct measurement.

In this section we shall relax the assumption, implicit in the von Neumann - Morgenstern setting, that the objective outcomes and subjective consequences coincide, and replace it with the weaker requirement that there exists a set of mutually exclusive, jointly exhaustive subjective consequences with respect to which the decision-maker evaluates each uncertain prospect. As we shall see such weaker condition, together with the usual assumptions of completeness and transitivity of preferences, and the Archimedean and Independence conditions from the von Neumann-Morgenstern treatment, jointly characterize representability in terms of a projective generalization of the expected utility functional.

Let $B := \{b^1, ..., b^n\}$ represent a set of n mutually exclusive, jointly exhaustive subjective consequences, identified with an orthonormal basis in $\mathbb{R}^n$. For any lottery $x \in X$, its associated *risk profile* p_x with respect to B is defined by

$$p_x(b^i) = \langle x | b^i \rangle^2, i = 1, \dots, n$$

The risk profile of a lottery x returns the probabilities induced on the subjective consequences by playing lottery x. Observe that the risk profile with respect to the natural basis simply returns the probabilities of the lottery outcomes. Hence, in the present setting a lottery is identified both in terms of objective outcomes and subjective consequences. Since the position of the subjective basis relative to the natural basis can be arbitrary, a lottery can exhibit any combination of risk profiles on outcomes and consequences.

Axiom 3 *(Born's Rule) There exists an orthonormal basis $Z := \{z^1, ..., z^n\}$ such that any two lotteries $x, y \in X$ are indifferent whenever their risk profiles with respect to Z coincide.*

Axiom 3 requires that there exists a set of n mutually exclusive, jointly exhaustive subjective consequences such that any two lotteries are only evaluated based on their risk profiles, *i.e.*, on the probabilities they induce on those subjective consequences. In the von Neumann - Morgenstern treatment, Axiom 3 is tacitly assumed to hold with respect to the natural basis in $\mathbb{R}^n$. This implicit assumption amounts to the requirement that lotteries are only evaluated based on the probabilities they induce on the objective lottery outcomes.

While the probabilities with respect to the natural basis represent the relevant dimensions of risk as perceived by the modeler or an external observer (that is, the risk associated with the occurrence of the objective outcomes), the *preferred basis* postulated in Axiom 3 is allowed to vary across different decision-makers, capturing the idea that the subjectively relevant dimensions of risk (that is, those pertaining to the actual subjective consequences) may be perceived differently by different subjects. One case in which the relevant dimensions of risk may differ across subjects is in the presence of portfolio effects, when the lottery outcomes are correlated with the portfolio outcomes. Such effects are very difficult to exclude or control for in experimental settings. For instance, the very fact of proposing to a subject the Allais or Ellsberg games described above generates an expectation of gain: an invitation to play the game can be effectively regarded as a risky security which is donated to the subject, and whose subjective returns are obviously correlated, but do not necessarily coincide, with the monetary outcomes of the experiment. Even in such simple contexts, significant hedging behavior cannot be in principle ruled out.

The preferred basis captures which, among all possible ways to partition the relevant uncertainty into a set of mutually exclusive events or outcomes, leads to a set of payoff-relevant, mutually orthogonal lotteries from which preferences on all other lotteries can be assigned in a linear fashion. In particular, the basis elements must span the whole range of preferences. For instance, in case a decision-maker is indifferent between two outcomes, but strictly prefers to receive them with equal frequency, preferences on the two outcomes do not span all the relevant range; therefore, this preference pattern cannot be captured in the natural basis, and hence in von Neumann - Morgenstern expected utility.

For now, and only for simplicity, we assume that subjective consequences and objective outcomes have the same cardinality; we relax this assumption later on, in the subjective formulation.

Once an orthonormal basis Z is given, each lottery x can be associated with a function $p_x : Z \to [0, 1]$, such that $p_x(z^i) = \langle x|z^i \rangle^2$ for all $z^i \in Z$. Let B be the set of all such risk profiles p_x, for $x \in X$, and let $\succsim$ be the complete and transitive preference ordering induced on $B \times B$ by preferences on the underlying lotteries.

Note that a convex combination $\alpha p_x + (1-\alpha)p_y$, where p_x and p_y are risk profiles, is still a well-defined risk profile. We interpret this type of mixing as objective, while subjective mixing will be later on captured by subjective probability over the underlying (Anscombe-Aumann) states. We postulate the following two axioms, which mirror those in the von Neumann - Morgenstern treatment.

Axiom 4 *(Archimedean) For all $x, y, t \in X$ with $p(x) \succ p(y) \succ p(t)$, there exist $\alpha, \beta \in (0, 1)$ such that $\alpha p(x) + (1 - \alpha)p(t) \succ p(y) \succ \beta p(x) + (1 - \beta)p(t)$.*

Axiom 5 *(Independence) For all $x, y, t \in X$, $p_x \succsim p_y$ if, and only if, $\alpha p_x + (1 - \alpha)p_t \succsim a p_y + (1 - \alpha)p_t$ for all $\alpha \in [0, 1]$.*

Some observations are in order at this point. First, note that the two axioms above impose conditions solely on risk profiles, and not on the underlying lotteries. This seems appropriate, as the decision-maker is not ultimately concerned with the risk associated to the objective outcomes, but only with the risk induced on the relevant subjective consequences. It is also worth noting that, while risk profiles are now defined with respect to subjective consequences, rather than objective outcomes as in the von Neumann - Morgenstern treatment, in our setting they are still interpreted as objective probability functions.

Theorem 2 *Axioms 3-5 are jointly equivalent to the existence of a symmetric matrix U such that $u(x) := x'Ux$ for all $x \in X$ represents $\succsim$.*

Proof. Assume that Axiom 3 holds with respect to a given orthonormal basis $\{z^1, ..., z^n\}$. By the von Neumann - Morgenstern result (which applies to any convex mixture set, such as B) Axioms 4 and 5 are jointly equivalent to the existence of a functional u which represents the ordering and is linear in p, i.e.

$$u(x) = \sum_{i=1}^{n} u(z^i)p_x(z^i) = \sum_{i=1}^{n} u(z^i) \langle x|z^i \rangle^2 ,$$

where the second equality is by definition of p as the squared inner product with respect to the preferred basis. The above can be equivalently written, in matrix form, as

$$u(x) = x'P'DPx = x'Ux,$$

where D is the diagonal matrix with the payoffs on the main diagonal, P is the projection matrix defined by $(P_{i,.})' = z^i$, and $U := P'DP$ is symmetric. Conversely, by the Spectral Decomposition theorem, for any symmetric matrix U there exist a diagonal matrix D and a projection matrix P such that $U = P'DP$, and hence

$$x'Ux = x'P'DPx$$

for all $x \in X$. But this is just expected utility with respect to the orthonormal basis defined by P. Hence, the three axioms are jointly equivalent to the existence of a symmetric matrix U such that $u(x) := x'Ux$ represents the preference ordering. QED

6 Subjective Formulation

The following formulation extends the representation to situations of subjective uncertainty. First, we introduce the following setup and notation.

$\mathcal{S}$ is a finite set of states of Nature.

$\langle .|. \rangle$ denotes the usual inner product in Euclidean space.

Ω is the natural basis in $\mathbb{R}^n$, identified with a finite set $\{\omega^1, \dots, \omega^n\}$ of lottery outcomes (prizes).

Z is an orthonormal basis in $\mathbb{R}^m$, with $m \geq n$, identified with a finite set of subjective consequences $\{z^1, \dots, z^m\}$. V is an arbitrary $(m \times n)$ matrix chosen so that, for all ω^i in Ω, $V\omega^i$ is a unit vector in $\mathbb{R}^m$. Observe that V is always well defined as long as $m \geq n$. When $m = n$, we conventionally set $V \equiv I$, where I is the $n \times n$ identity matrix.

Lotteries correspond to L^2 unit vectors $x \in \mathbb{R}^n_+$; X is the set of all lotteries.

Since Ω is the natural basis, $\langle \omega^i|x \rangle^2 = x_i^2$; this quantity is interpreted as $p(\omega^i|x)$.

The quantity $\langle z^j|Vx \rangle^2$ is interpreted as $p(z^j|x)$, the conditional probability of subjective consequence z^j given lottery x. In particular, $\langle z^j|V\omega^i \rangle^2$ is interpreted as $p(z^j|\omega^i)$, the conditional probability of subjective consequence z^j given the degenerate lottery which returns objective outcome ω^i for sure.

Once the subjective consequences z^j are specified, for any lottery x one can readily compute $p(\omega^i|x) = x_i^2$ and $p(z^j|x) = \langle z^j|Vx\rangle^2$. Moreover, given the latter probabilistic constraints, one can readily identify a lottery x and an orthonormal basis Z which jointly satisfy them. Hence, in the above construction lotteries are identified with respect to two different frames of reference: objective lottery outcomes, and subjective consequences.

Observe that $p(z^j|x)$ generally differs from the probability of z^j given x computed according to the law of total probability, which is given by $\sum_i p(\omega^i|x)p(z^j|\omega^i) = \sum_i x_i^2 \langle z^j|V\omega^i\rangle^2$. To get a sense of why and how the law of total probability may fail, let us consider a decision-maker who really hates to lose whenever the probability of winning is high (more so than when the probability of winning is low), and loves to win when the probability of losing is high (even more so than when the latter probability is low). Clearly, in such case the probabilities of objective outcomes such as winning and losing are directly involved in the description of subjectively relevant consequences such as "I won (or lost) against all odds". Such dependency introduces an element of interference between lotteries and consequences that cannot be easily accounted for in the classical decision-theoretic setting, which presumes state independence.

An act is identified with a function $f : \mathcal{S} \to X$. H is the set of all acts.

$\Delta(X)$ is the (nonempty, closed and convex) set of all probability functions on Z induced by lotteries in X.

M is the set of all vectors $(p_s)_{s\in S}$, with $p_s \in \Delta(X)$.

For each $f \in H$ a corresponding risk profile $p^f \in M$ is defined, for all $s \in \mathcal{S}$ and all $z^j \in Z$, by $p_s^f(z^j) := \langle z^j|Vf_s\rangle^2$.

As customary, we assume that the decision-maker's preferences are characterized by a rational (*i.e.*, complete and transitive) preference ordering $\succsim$ on acts.

Next, we proceed with the following assumptions, which mirror those in Anscombe and Aumann (1963).

Axiom 6 *(Projective) There exists a finite orthonormal basis $Z := \{z_1, ..., z_m\}$, with $m \geq n$, such that any two acts $f, g \in H$ are indifferent if $p^f = p^g$.*

In Anscombe and Aumann's setting, the above axiom is implicitly assumed to hold with $Z \equiv \Omega$. Because of Axiom 6, preferences on acts can be equivalently expressed as preferences on risk profiles. For all $p^f, p^g \in M$, we stipulate that $p^f \succsim p^g$ if and only if $f \succsim g$.

Axiom 7 *(Archimedean) If $p^f, p^g, p^h \in M$ are such that $p^f \succ p^g \succ p^h$, then there exist $a, b \in (0, 1)$ such that $ap^f + (1-a)p^h \succ p^g \succ bp^f + (1-b)p^h$.*

Axiom 8 *(Independence) For all $p^f, p^g, p^h \in M$, and for all $a \in (0, 1]$, $p^f \succ p^g$ if and only if $ap^f + (1-a)p^h \succ ap^g + (1-a)p^h$.*

Axiom 9 *(Non-degeneracy) There exist $p^f, p^g \in M$ such that $p^f \succ p^g$.*

Axiom 10 *(State independence) Let $s, t \in \mathcal{S}$ be non-null states, and let $p, q \in \Delta(X)$. Then, for any $p^f \in M$,*

$$(p_1^f, ..., p_{s-1}^f, p, p_{s+1}^f, ..., p_n^f) \succ (p_1^f, ..., p_{s-1}^f, q, p_{s+1}^f, ..., p_n^f)$$

if, and only if,

$$(p_1^f, ..., p_{t-1}^f, p, p_{t+1}^f, ..., p_n^f) \succ (p_1^f, ..., p_{t-1}^f, q, p_{t+1}^f, ..., p_n^f).$$

Theorem 3 *(Anscombe and Aumann) The preference relation $\succsim$ fulfills Axioms $6-10$ if and only if there is a unique probability measure π on $\mathcal{S}$ and a non-constant function $u : Z \to \mathbb{R}$ (unique up to positive affine rescaling) such that, for any $f, g \in H$, $f \succsim g$ if, and only if,*

$$\sum_{s\in\mathcal{S}} \pi(s) \sum_{z^i\in Z} p_s^f(z^i)u(z^i) \geq \sum_{s\in\mathcal{S}} \pi(s) \sum_{z^i\in Z} p_s^g(z^i)u(z^i).$$

The following result provides a projective generalization of the Anscombe and Aumann theorem.

Theorem 4 *The preference relation $\succsim$ fulfills Axioms $6-10$ if and only if there is a unique probability measure π on $\mathcal{S}$ and a symmetric $(n \times n)$ matrix U with distinct eigenvalues such that, for any $f, g \in H$, $f \succsim g$ if, and only if,*

$$\sum_{s\in\mathcal{S}} \pi(s)f_s'Uf_s \geq \sum_{s\in\mathcal{S}} \pi(s)g_s'Ug_s.$$

Proof. Let D be the $(m \times m)$ diagonal matrix defined by $D_{i,i} = u(z^i)$, and let P be the projection matrix defined by $(P_{i,.})' = z^i$. If Axioms 6-10 hold, we know from Theorem 3 that the preference ordering has an expected utility representation. Observe that, since $(PVf_s)_i^2 = \langle z^i|Vf_s\rangle^2 = p_s^f(z^i)$, the expected utility of any act f can be written as

$$\sum_{s\in\mathcal{S}} \pi(s) \sum_{z^i\in Z} p_s^f(z^i)u(z^i) = \sum_{s\in\mathcal{S}} \pi(s)f_s'Uf_s,$$

where $U := V'P'DPV$ is a $(n \times n)$ symmetric matrix. Furthermore, observe that at least two eigenvalues of U must be distinct, otherwise $u(f) = c$ for all $f \in H$, contradicting Axiom 9.

Conversely, let U be a symmetric $(n \times n)$ matrix with distinct eigenvalues. By the Spectral Decomposition theorem there exist a projection matrix P and a diagonal matrix D such that $U = P'DP$, and therefore

$$\sum_{s \in \mathcal{S}} \pi(s) f_s' U f_s = \sum_{s \in \mathcal{S}} \pi(s) f_s' P' D P f_s.$$

Observe that the right-hand side is the expected utility of f with respect to the n-dimensional orthonormal basis Z defined by $z^i = P'_{i,\cdot}$, and with $u(z^i) = D_{i,i}$. Since the diagonal elements of D are the eigenvalues of U, if the latter has distinct eigenvalues then $u(z^i)$ is non-constant, and hence by Theorem 3 Axioms 6-10 must hold. QED

7 Properties of the Representation

Our representation generalizes the Anscombe-Aumann expected utility framework in three directions. First, subjective uncertainty and risk are treated as distinct notions. Specifically, let us say that an act is $pure$ or $certain$ if it returns the same objective lottery in all states, and $mixed$ or $uncertain$ otherwise. While pure acts are naturally associated with risky decisions, in which the relevant frequencies are all known, mixed acts correspond to uncertain decisions, in which the decision-maker only has a subjective assessment of the true frequencies involved. Second, as we shall see, within this class of preferences both Allais' and Ellsberg's paradoxes are accommodated. Finally, the construction easily extends to the complex unit sphere, and hence to wave functions, provided that $\langle x|y\rangle^2$ is replaced by $|\langle x|y\rangle|^2$ in the definition of p, in which case Theorems 2 and 4 hold with respect to a Hermitian (rather than symmetric) payoff matrix U, and the result also provides axiomatic foundations for decisions involving quantum uncertainty.

In the context of Theorem 2, for any two distinct outcomes ω^i and ω^j let $e_{i,j}$ be the objective lottery returning each of the two outcomes with equal frequency. Observe that

$$U_{ij} = u(e_{i,j}) - (\frac{1}{2}u(\omega^i) + \frac{1}{2}u(\omega^j)).$$

It follows that the off-diagonal entry U_{ij} in the payoff matrix can be interpreted as the discount, or premium, attached to a symmetric, objective lottery over the two outcomes with respect its expected utility base-line, and hence as a measure of preference for risk versus uncertainty along the specific dimension involving outcomes ω^i and ω^j. Let us say that a decision-maker is *averse to uncertainty* if she always weakly prefers $e_{i,j}$ to an equal subjective chance of ω^i or ω^j. Then a decision-maker is averse to uncertainty if and only if U is a Metzler matrix, *i.e.*, has non-negative off-diagonal elements.

Observe that the functional in Theorem 2 is quadratic in x, but linear in p. If U is diagonal, then its eigenvalues coincide with the diagonal elements. In von Neumann - Morgenstern expected utility, those eigenvalues contain all the relevant information about the decision-maker's risk attitudes. By contrast, in our setting risk attitudes with respect to objective outcomes are captured by both the diagonal and non-diagonal elements of U. In turn, U is jointly characterized by a projection matrix P, and a diagonal matrix D with the eigenvalues of U on the main diagonal. While P identifies a set of subjective consequences, D captures the decision maker's risk attitudes with respect to those consequences.

Compared to existing generalizations of expected utility which avoid the Allais or Ellsberg paradoxes, such as the ones in Machina (1982), Schmeidler (1989), or Chew, Epstein and Segal (1991), among others, projective expected utility enjoys several advantages. Specifically, the representation is linear in the probabilities of states and consequences, hence remaining quite tractable, but can be nonlinear in the probabilities of the objective outcomes, hence allowing for portfolio effects. The axioms used to obtain the representation closely mirror those introduced by Anscombe and Aumann (1963), which are widely regarded as appealing. Moreover, as shown in sections 8 and 9, a compatible specification of the payoff matrix avoids both Allais' and Ellsberg's paradox. Finally, in case the event space is non-classical, the result also provides foundations for decisions involving quantum uncertainty. An early paper combining ideas from quantum mechanics and decision theory is Deutsch (1999), which provides a decision-theoretic foundation for quantum probability in the special case of purely monetary outcomes. Our generalization of expected utility is closest to the one in Gyntelberg and Hansen (2004), which is obtained in a Savage context by postulating a non-classical (that is, non-Boolean) structure for the relevant events. By contrast, our representation is obtained in an Anscombe-Aumann context, and does not impose specific requirements on the nature of the relevant uncertainty. In particular, in our setting the dominant paradoxes can be resolved even if the relevant uncertainty is of completely classical nature.

8 Example: Objective Uncertainty

Figure 1 below presents several examples of indifference maps on pure lotteries which can be obtained within our class of preferences for different choices of U.

The indifference maps refer to decision problems with

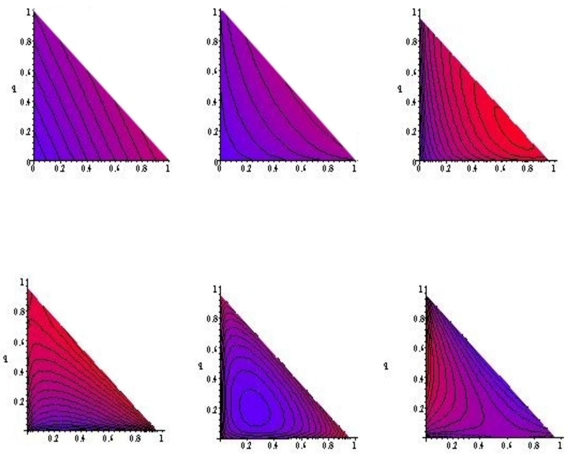

Figure 1: Examples of indifference maps on the probability triangle, with payoff matrices given, respectively, by

3	0	0
0	2	0
0	0	1

3	1	0
1	2	0
0	0	1

3	1	1
1	2	0
1	0	1

3	1	-2
1	2	1
-2	1	1

3	-1	-2
-1	2	-1
-2	-1	1

3	-4	0
-4	2	5
0	5	1

three objective outcomes, which correspond to the vertices of the triangle. All other points in the triangle represent probability distributions over those three alternatives, and the iso-level curves represent loci of equal utility (indifference curves). Darker shades correspond to lower utility. The first pattern (parallel straight lines) characterizes von Neumann - Morgenstern expected utility. Within our class of representations, it corresponds to the special case of a diagonal payoff matrix U. All other patterns are impossible within von Neumann - Morgenstern expected utility. Observe that, even though the payoff matrix is an object of relatively limited algebraic complexity, the indifference curves can take a variety of different shapes: in particular, they do not need to be convex, or concave. Positive off-diagonal payoffs imply a premium for probabilistic mixtures over the corresponding alternatives, while negative payoffs imply a discount relative to the expected utility baseline. Clearly, since the indifference maps are generated by a limited number of parameters (the entries in the payoff matrix), the type and variety of preference patterns predicted by the model is also limited, and this in turn offers a basis for the empirical testability of the theory.

The representation is sufficiently general to accommodate Allais' paradox from section 3. In the context of the example in section 3, let $\{\omega^1, \omega^2, \omega^3\}$ be the outcomes in which 4000, 3000, and 0 dollars are won, respectively. To accommodate Allais' paradox, assume that the non-diagonal elements of the payoff matrix increase with the difference between the corresponding diagonal payoffs. In the example below, the non-diagonal elements are taken to be proportional to the fourth power of the difference.

$$U = \begin{array}{cccc} & \omega^1 & \omega^2 & \omega^3 \\ \omega^1 & 1.1 & 0.00001 & 0.14641 \\ \omega^2 & 0.00001 & 1 & 0.1 \\ \omega^3 & 0.14641 & 0.1 & 0 \end{array}$$

The above formulation of the payoff matrix implies that, whenever the stakes involved are similar, the corresponding prospects are evaluated approximately at their (von Neumann-Morgenstern) expected utility values. By contrast, whenever the stakes are significantly different, the divergence from expected utility is also significant. Let the four lotteries A, B, C, D be defined, respectively, as the following unit vectors in $\mathbb{R}^3_+$: $a := (\sqrt{0.2}, 0, \sqrt{0.8})'$; $b := (0, \sqrt{0.25}, \sqrt{0.75})'$; $c := (\sqrt{0.8}, 0, \sqrt{0.2})'$; $d := (0, 1, 0)'$. Then lottery A is preferred to B, while D is preferred to C, as

$$u(a) = a'Ua = 0.33713,$$
$$u(b) = b'Ub = 0.3366,$$
$$u(c) = c'Uc = 0.99713,$$
$$u(d) = d'Ud = 1.$$

9 Example: Subjective Uncertainty

In the Ellsberg puzzle, suppose that either all the non-red balls are green (*i.e.*, 100 red, 200 green, 0 blue), or they are all blue (100 red, 0 green, 200 blue), with equal subjective probability. Further, suppose that there are just two objective outcomes, Win and Lose. Then the following specification of the payoff matrix accommodates the paradox.

$$U = \begin{array}{ccc} & \text{Win} & \text{Lose} \\ \text{Win} & 1 & \alpha \\ \text{Lose} & \alpha & 0 \end{array}$$

As we shall see, if $\alpha = 0$ we are in the expected utility case, where the decision-maker is indifferent between risk and uncertainty; when $\alpha > 0$, risk is preferred to uncertainty; and when $\alpha < 0$, the decision-maker prefers uncertainty to risk.

In fact, let $\{Urn1, Urn2\}$ be the set of possible states of nature, with uniform subjective probability, and let

$$r := (\sqrt{1/3}, \sqrt{2/3})',$$
$$\overline{r} := (\sqrt{2/3}, \sqrt{1/3})'$$

be the lotteries associated to pure acts R and $\overline{R}$, respectively. Furthermore, let

$w := (1, 0)'$,

$l := (0, 1)'$

be the lotteries corresponding to a sure win (W) and a sure loss (L), respectively.

The mixed acts G and $\overline{G}$ have projective expected utilities given by

$$u(G) = p(Urn1)u(\overline{R}) + p(Urn2)u(L),$$
$$u(\bar{G}) = p(Urn1)u(R) + p(Urn2)u(W).$$

One also has that

$$u(W) = w'Uw = 1,$$
$$u(L) = l'Ul = 0,$$
$$u(R) = r'Ur = 1/3 + \alpha\sqrt{8}/3,$$
$$u(\overline{R}) = \overline{r}'U\overline{r} = 2/3 + \alpha\sqrt{8}/3,$$

and therefore

$$u(G) = 1/3 + \alpha\sqrt{2}/3$$
$$u(\overline{G}) = 2/3 + \alpha\sqrt{2}/3.$$

It follows that, whenever $\alpha > 0$, R is preferred to G and $\overline{R}$ to $\bar{G}$, so the paradox is accommodated. When $\alpha < 0$, the opposite pattern emerges: G is preferred to R and $\bar{G}$ to $\overline{R}$. Finally, when $\alpha = 0$ the decision-maker is indifferent between R and G, and between $\overline{R}$ and $\bar{G}$.

10 Games with PEU preferences

Within the class of preferences characterized by Theorem 4, is it still true that every finite game has a Nash equilibrium? If the payoff matrix U is diagonal we are in the classical case, so we know that any finite game has an equilibrium, which moreover only involves objective risk (in our terms, this type of equilibrium should be referred to as "pure", as it involves no subjective uncertainty). For the general case, consider that $u(f)$ is still continuous and linear with respect to the subjective beliefs π, while possibly nonlinear (but still polynomial) with respect to risk. As we shall see, any finite game has an equilibrium even within this larger class of preferences, although the equilibrium may not be pure (in our sense): in general, an equilibrium will rest on a combination of objective randomization and subjective uncertainty about the players' decisions.

A finite, strategic-form game with PEU preferences is a n-tuple $G := (I, (A^i)_{i\in I}, (U^i)_{i\in I})$, where I is a finite set composed of k players, A^i represents the set of feasible actions for player i, and U^i is player i's payoff matrix on outcomes. An outcome is a complete assignment of actions $(a^1, ..., a^k)$, identified with the natural basis in R^d, where $d := \prod_i |A^i|$. An act of player i is a function f^i from states S^i to lotteries on A^i. Let $\mathcal{H}^i$ be the set of feasible acts for player i. We assume that $\mathcal{H}^i$ is a compact set which is also convex with respect to both objective and subjective mixing, in the following sense.

Assumption (Convexity): For all $f^i, g^i \in \mathcal{H}^i$, and for all $a \in [0, 1]$, there exist $h^i, l^i \in \mathcal{H}^i$ such that

(i) $p_s^{h^i} = ap_s^{f^i} + (1-a)p_s^{g^i}$ for all s (objective mixture)

(ii) $u(l^i) = au(f^i) + (1-a)u(g^i)$ (subjective mixture).

A profile of acts is a unitary vector $f := (f^1, ..., f^k)$, and the utility of a profile f for player i is $u^i(f) = E_\pi[f'U^i f]$. An equilibrium is a profile f^* such that $u^i(f_i^*, f_{-i}^*) \geq u^i(f_i, f_{-i}^*)$ for all f^i.

Theorem: Any finite, strategic-form game with PEU preferences has an equilibrium.

Proof: First, let $b^i(f^{-i}) := \{f^i : [u^i(f^i, f^{-i}) \geq u^i(g^i, f^{-i})](\forall g^i \in \mathcal{H}^i)\}$. Next, let $\mathcal{H} := \times_i \mathcal{H}^i$, and let $b : \mathcal{H} \to \mathcal{H}$ be the best response correspondence, defined by $b(f) = \times_i b^i(f^i)$. Observe that b is a correspondence from a nonempty, convex, and compact set $\mathcal{H}$ to itself. In addition, b is a nonempty- and convex-valued, upper hemicontinuous correspondence. It follows that the conditions of Kakutani's fixed point theorem are satisfied, and hence the best response correspondence b has a fixed point: a profile of acts f^* such that $f^* \in b(f^*)$. The acts at this fixed point constitute an equilibrium since by construction $f^{*i} \in b^i(f^{*-i})$ for all i.QED

11 Games with PEU preferences: Examples

Consider the following game:

$Pl.1, Pl.2$	a^2	b^2
a^1	$1.1, 0$	$1, 1.1$
b^1	$0, 1.1$	$1.1, 1$

If the agents maximize vN-M expected utility, the unique mixed strategy equilibrium takes the following form:

$(p, q) = (1/12, 1/12) \approx (0.08333, 0.08333)$

where p and q are the probabilities of playing a^1 and a^2, respectively. We shall call this equilibrium EU equilibrium. If the players have the preferences we introduced to explain Allais paradox, what are the consequences on the equilibrium strategies of the players? In particular, does the equilibrium differ in the case of Allais agents?

Consider the following game:

$Pl.1, Pl.2$	a^2	b^2
a^1	ω^1, ω^3	ω^2, ω^1
b^1	ω^3, ω^1	ω^1, ω^2

where the payoff matrix of both players is given by U in the Allais example. Observe that the payoff matrix U describes a decision-maker who is strictly uncertainty averse, and hence if all players have payoffs matrices given by U any equilibrium only involves pure acts. In this case, the unique equilibrium is given by

$$(p, q) = (0.17632, 0.17632).$$

Compared to the case of EU preferences, in the game with PEU preferences a^1 and a^2 are played more often. Starting from the EU equilibrium, observe that strategy a^2 becomes more attractive for P2. To re-establish equilibrium a^1 must be played more often, but then to keep P1 indifferent between her two strategies a^2 must also be played more often.

Next, consider the following very simple version of the Centipede game: P1 either quits (outcome; ω^2, ω^3) or passes, in which case P2 either quits (outcome; ω^3, ω^1), or passes (outcome; ω^1, ω^2). The payoff matrices of the two players are still given by U.

Let p denote the probability that P1 passes, and let q be the probability that P2 passes given that P1 passed. Then the PEU for P1 and P2 are given, respectively, by

$$((1-p)(1-q), p, (1-p)q)U((1-p)(1-q), p, (1-p)q)'$$

$$((1-p)q, (1-p)(1-q), p)U((1-p)q, (1-p)(1-q), p)'.$$

In a subgame-perfect equilibrium, player 2 chooses q so that it maximizes $(\sqrt{q}, \sqrt{(1-q)}, 0)U(\sqrt{q}, \sqrt{(1-q)}, 0)'$, and hence will choose $q = 1$. Given that $q = 1$, P1 will choose p to maximize $(0, \sqrt{p}, \sqrt{(1-p)})U(0, \sqrt{p}, \sqrt{(1-p)})'$, and hence in equilibrium $p = 0.99029$. It follows that the unique subgame-perfect equilibrium in this game involves a small probability of continuation. Clearly, in a longer version of the centipede these continuation probability would be amplified.

12 References

Allais, M. 1953. Le Comportement de l'Homme Rationnel devant le Risque: Critique des postulats et axiomes de l'École Americaine. *Econometrica* 21: 503-546.

Anscombe, F.J. and Aumann, R.J 1963. A Definition of Subjective Probability. *The Annals of Mathematical Statistics* 34: 199-205.

Bernoulli, D. 1738. Specimen theoriae novae de mensura sortis. In Commentarii Academiae Scientiarum Imperialis Petropolitanae, 5: 175-192. Translated as "Exposition of a new theory of the measurement of risk". 1954, *Econometrica* 22: 123-136.

Chew, S.H., L.G. Epstein and U. Segal, 1991, Mixture symmetry and quadratic utility, *Econometrica* 59, 139-164.

Deutsch, D. 1999. Quantum Theory of Probability and Decisions. In *Proceedings of the Royal Society* A455: 3129-3197.

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

Gyntelberg, J. and Hansen, F. 2004. Expected utility theory with small worlds, Department of Economics Discussion Paper 04-20, University of Copenhagen.
http://ideas.repec.org/p/kud/kuiedp/0420.html

Machina, M. J. 1982. Expected utility analysis without the independence axiom, *Econometrica* 50: 277-323.

von Neumann, J. 1932. *Mathematische Grundlagen der Quantenmechanic*. Berlin, Springer. Translated as *Mathematical foundations of quantum mechanics*, 1955, Princeton University Press, Princeton NJ.

von Neumann, J., and Morgenstern, O. 1944. *Theory of Games and Economic Behavior*. Princeton University Press, Princeton NJ.

Rabin, M. and Thaler, R. H. 2001. Anomalies: Risk Aversion. *Journal of Economic Perspectives* 15: 219-232.

Savage, L. J. 1954. *The Foundations of Statistics*. Wiley, New York, NY.

Schmeidler, D. 1989. Subjective probability and expected utility without additivity. *Econometrica*, 57: 571-587.

Tversky, A. 1975. A critique of expected utility theory: Descriptive and normative considerations. *Erkenntnis*, 9: 163-173.