Notes on Expected Utility and Uncertainty Aversion

Pierpaolo Battigalli
Bocconi University and IGIER

Abstract

These notes complement Chapter 3 of *Game Theory, Analysis of Strategic Thinking.* That chapter only covers subjective expected utility maximization. Here we also cover the model of maxmin expected utility and the so called "smooth ambiguity" model. The decision theoretic axiomatizations of these models are not covered, we just look at functional-form representations. We stress the geometrical interpretation and present some results at the advanced-undergraduate level, without using a rigid definition-theorem-proof format.

1 Setting

We consider a problem of decision under uncertainty. This problem may be a game between the Decision Maker (DM), a player on whom we focus our attention, and another player whom we treat as "Nature." We compare some decision criteria that have been axiomatized in decision theory in terms of preferences over acts (see the surveys by Gilboa and Marinacci [7] and Marinacci [10]). But our analysis does not concern the axiomatization (for an axiomatization close to the formalism of these notes see Battigalli *et al.* [6]).

1.1 Primitives

All the sets we consider below are *finite:*

- **Actions** $a \in A$
- **States** $s \in S$
- **Outcomes** (or consequences) $y \in Y$
- **Outcome function** $g : A \times S \to Y$
- Von Neumann Morgenstern (VNM) **utility function**[1] $v : Y \to \mathbb{R}$

From $g : A \times S \to Y$ and $v : Y \to \mathbb{R}$ we derive the **payoff function**

$$u = v \circ g : A \times S \to \mathbb{R}.$$

[1]See the Appendix of [12].

1.2 Objective probabilities, lotteries and state-contingent lotteries

We call **probability distribution** on a finite set X an element of the unit simplex in $\mathbb{R}^X$, i.e., an element of the set

$$\Delta(X) := \left\{ \xi \in \mathbb{R}_+^X : \sum_{x \in X} \xi(x) = 1 \right\}.$$

The support of a probability distribution $\xi \in \Delta(X)$ is the set of points in X with positive probability:

$$\mathrm{supp}\xi := \{x \in X : \xi(x) > 0\}.$$

1.2.1 Risk

A decision problem under **risk** is characterized by a known *objective* probability distribution $\sigma \in \Delta(S)$. Whether a probability is objective or subjective is a matter of interpretation.[2] We use the symbol σ to stress the objective interpretation. Think of σ as the composition of an urn, where S is the set of colors, e.g., the urn contains balls of two colors, lilac (ℓ) and red (r), and $\sigma(r)$ is the fraction of red balls. In a decision problem under risk, DM knows the composition of the urn, hence the objective probability of drawing a ball of each given color. Therefore DM also knows that each action $a \in A$ induces a corresponding objective probability distribution over consequences, called **lottery**:

$$\begin{array}{rcl} a & \longmapsto & \lambda_{a,\sigma} \in \Delta(Y), \\ \lambda_{a,\sigma}(y) & = & \displaystyle\sum_{s: g(a,s)=y} \sigma(s). \end{array}$$

Since we want to study the incentive to randomize, we are going to consider a larger decision problem whereby DM has a rich *commitment technology* allowing him to pick an objective random device and let his actual action be chosen by this device. For example, DM can spin a roulette wheel and let a machine (or trustworthy agent) pick a if the realization is an even number and b if it is an odd number, or alternatively he can let a machine pick a if the realization is a number $x \leq 10$ and b if $x > 10$. The resulting probabilities of actions are denoted $\alpha(a)$, the resulting mixed action is denoted $\alpha = (\alpha(a))_{a \in A} \in \Delta(A)$. By choosing the random device and the rule that maps realizations to actions, the DM determines a mixed action $\alpha \in \Delta(A)$. Note, it is important that DM irreversibly commits to let the random device select the pure action, for suppose that the DM could change the decision after the realization of the random device (e.g., the roulette wheel), if the pure action selected by device is not deemed optimal by DM, the he would switch ex post to another action. This cannot happen when DM is a subjective expected utility maximizer, but it may happen according to other decision models considered below.

In general, the available commitment technology allows to implement some, but possibly not all the mixed actions; thus DM must choose within a given feasible subset $\mathcal{A} \subseteq \Delta(A)$

[2] Some probability theorists even deny the existence of objective probabilities.

of mixed actions. But here, with the exception of Section 6, we assume for simplicity that the commitment technology is rich enough to span the whole space of mixed actions, i.e., $\mathcal{A} = \Delta(A)$.

We will see that randomizations are redundant for a subjective expected utility maximizer, but not for other types of decision makers. Be as it may, given the objective distribution σ over states, each mixed action α induces a lottery $\lambda_{\alpha,\sigma} \in \Delta(Y)$, which is the convex combination of the lotteries $\lambda_{a,\sigma}$ ($a \in \mathrm{supp}\alpha$) with weights given by the probabilities of pure actions:

$$\lambda_{\alpha,\sigma} = \sum_{a \in A} \alpha(a) \lambda_{a,\sigma},$$

that is, for each consequence $y \in Y$,

$$\lambda_{\alpha,\sigma}(y) = \sum_{a \in A} \alpha(a) \lambda_{a,\sigma}(y) = \sum_{a \in A} \sum_{s : g(a,s) = y} \alpha(a) \sigma(s).$$

1.2.2 Uncertainty

If DM does not know the objective distribution (urn composition) $\sigma \in \Delta(S)$, we have a problem of decision under **uncertainty.** All we can say in this case is that, for each $s \in S$, a mixed action $\alpha \in \Delta(A)$ induces the corresponding lottery

$$\begin{array}{rcl} \lambda_{\alpha,s} & \in & \Delta(Y) \\ y & \longmapsto & \lambda_{\alpha,s}(y) = \sum_{a : g(a,s) = y} \alpha(a). \end{array}$$

In words, $\lambda_{\alpha,s}$ is the lottery that picks consequence c with the total probability of actions that yield c in state s. Thus, a mixed action does not induce a single lottery, but a profile of lotteries, one for each state s:

$$\lambda_\alpha = (\lambda_{\alpha,s})_{s \in S} \in [\Delta(Y)]^S.$$

Such profiles are called **Anscombe–Aumann acts** in decision theory (see [1]). Thus, the set of mixed actions induces a corresponding subset of the set of Anscombe–Aumann acts:

$$\{\lambda \in [\Delta(Y)]^S : \exists \alpha \in \Delta(A), \lambda = \lambda_\alpha\} \subseteq [\Delta(Y)]^S.$$

2 Intermezzo: probability distributions, probability measures, pushforward

Another useful way to write the induced lottery $\lambda_{a,\sigma}$ involves the isomorphism between probability distributions and probability measures on a finite space, and the relationship between probability measures on different spaces, in our case S and Y.

2.1 Probability distributions and measures

First, note that for every *finite* set X there is an isomorphism between the unit simplex in $\mathbb{R}^X$ and the set of probability measures over 2^X (the power set of X)[3]

$$\left\{\mu \in \mathbb{R}_+^{2^X} : \mu(X) = 1, \forall E, F \in 2^X, E \cap F = \varnothing \Rightarrow \mu(E \cup F) = \mu(E) + \mu(F)\right\}.$$

The isomorphism is obvious: each measure μ (a function in $\mathbb{R}_+^{2^X}$) determines the distribution ξ_μ (a function in $\mathbb{R}_+^X$) where

$$\forall x \in X,\ \xi_\mu(x) = \mu(\{x\});$$

each distribution ξ determines the measure μ_ξ where

$$\forall E \in 2^X,\ \mu_\xi(E) = \sum_{x \in E} \xi(x).$$

It is easy to check that $\mu_{\xi_\mu} = \mu$ and $\xi_{\mu_\xi} = \xi$, and that convex combinations (also called mixtures) of distributions and measures are preserved under such invertible map: for all $\gamma \in \Delta(\{1, ..., n\})$, $\xi_1, ..., \xi_n$, $\mu_1, ..., \mu_n$,

$$\begin{aligned}
\sum_{k=1}^n \gamma_k \mu_{\xi_k} &= \mu_{\left(\sum_{k=1}^n \gamma_k \xi_k\right)}, \\
\sum_{k=1}^n \gamma_k \xi_{\mu_k} &= \xi_{\left(\sum_{k=1}^n \gamma_k \mu_k\right)}.
\end{aligned}$$

Given the isomorphism between the space of probability distributions and the space of probability measures, it makes sense denote both with the symbol $\Delta(X)$.

2.2 Pushforward

Take a map f between two sets,[4] say

$$f : X \to Y.$$

It is intuitive that each probability measure $\mu \in \Delta(X)$ induces, through f, a corresponding probability measure $\nu = \hat{f}(\mu)$ on Y. Such induced measure $\hat{f}(\mu)$ is called the **pushforward** of μ through f. Here is the formula:

$$\begin{aligned}
\hat{f}(\mu) &= \mu \circ f^{-1} \in \Delta(Y), \\
E &\longmapsto \hat{f}(\mu)(E) = \mu(f^{-1}(E)).
\end{aligned}$$

[3] More generally, a probability measure can be defined for every **measurable space** $(X, \mathcal{F})$, where $\mathcal{F} \subseteq 2^X$ is a collection of subsets, called sigma-algebra, or sigma-field, such that (i) $X \in \mathcal{F}$, and for every $E, E_1, E_2, ... \in 2^X$ (ii) $E \in \mathcal{F} \Rightarrow X \backslash E \in \mathcal{F}$, (iii) $E_1, E_2, ... \in \mathcal{F} \Rightarrow \cup_n E_n \in \mathcal{F}$. A **probability measure** on $(X, \mathcal{F})$ is a map $\mu : \mathcal{F} \to [0,1]$ such that (1) $\mu(X) = 1$, and (2) if $E_1, E_2, ... \in \mathcal{F}$ are pairwise disjoint then $\mu(\cup_n E_n) = \sum_n \mu(E_n)$. When X is finite it is customary to consider $\mathcal{F} = 2^X$ and we are back to the definition in the main text.

[4] We keep considering finite spaces: but what we say here applies to infinite measurable spaces as well, i.e., spaces endowed with a sigma-algebra of measurable events, like the interval $[0,1]$ with the collection of all Borel subsets of $[0,1]$.

In other words, for each subset $E \subseteq Y$, the induced probability of E is the summation of the probabilities of pre-images in X of points $y \in E$

$$\hat{f}(\mu)(E) = \sum_{x \in f^{-1}(E)} \mu(\{x\}).$$

2.3 Sections of functions and acts

Next, note that each action a determines a map from states to consequences, which is the **section g_a of function g at a**:

$$\begin{array}{rcl} g_a & : & S \to Y, \\ s & \longmapsto & g(a,s) \in Y. \end{array}$$

A map from states to consequences is called **Savage act** in decision theory (see Savage [11]). The set of pure actions induces a subset of the set of savage acts:

$$\{f \in Y^S : \exists a \in A, f = g_a\} \subseteq Y^S.$$

With this, for each objective distribution (or measure) $\sigma \in \Delta(S)$ we can write the lottery $\lambda_{a,\sigma}$ induced by pure action a as follows

$$\begin{array}{rcl} \lambda_{a,\sigma} & = & \hat{g}_a(\sigma) = \sigma \circ g_a^{-1} \in \Delta(Y) \\ E & \longmapsto & \lambda_{a,\sigma}(E) = \sigma(g_a^{-1}(E)). \end{array}$$

Similarly, each mixed action $\alpha \in \Delta(A)$ also induces a corresponding lottery in $\Delta(Y)$, which is the weighted sum of the lotteries induced by the pure actions $a \in \mathrm{supp}\alpha$:

$$\lambda_{\alpha,\sigma} = \sum_{a \in A} \alpha(a) \lambda_{a,\sigma},$$

that is, for each subset $E \subseteq Y$,

$$\lambda_{\alpha,\sigma}(E) = \sum_{a \in A} \alpha(a) \lambda_{a,\sigma}(E).$$

If no objective measure $\sigma \in \Delta(S)$ is given, we cannot derive a lottery in $\Delta(Y)$ from a pure or mixed action α. All we can say is that, for each $s \in S$, the mixed action α induces a lottery $\lambda_{\alpha,s} \in \Delta(Y)$ obtained as the pushfoward of α on Y through the section $g_s : A \to Y$

$$\begin{array}{rcl} \lambda_{\alpha,s} & = & \hat{g}_s(\alpha) = \alpha \circ g_s^{-1} \\ E & \longmapsto & \lambda_{\alpha,s}(E) = \alpha(g_s^{-1}(E)). \end{array}$$

This yields the profile of lotteries (Anscombe–Aumann act)

$$\lambda_\alpha = (\lambda_{\alpha,s})_{s \in S} \in [\Delta(Y)]^S.$$

3 Risk: objective expected utility

When we say that $v : Y \to \mathbb{R}$ is a **VNM utility function** we mean that in decision problems under risk DM ranks lotteries in $\Delta(Y)$ according to the corresponding expected value of v:

$$\lambda' \succsim \lambda'' \iff \sum_{y \in Y} \lambda'(y)v(y) \geq \sum_{y \in Y} \lambda''(y)v(y),$$

or, more briefly,[5]

$$\mathbb{E}_{\lambda'}[v] \geq \mathbb{E}_{\lambda''}[v].$$

Therefore, for a given objective probability measure (urn composition) σ, DM ranks actions $a \in A$ as follows:

$$a \succsim b \iff \mathbb{E}_{\lambda_{a,\sigma}}[v] \geq \mathbb{E}_{\lambda_{b,\sigma}}[v],$$

or, more explicitly,

$$a \succsim b \iff \sum_{y \in Y} \sigma(g_a^{-1}(y))v(y) \geq \sum_{y \in Y} \sigma(g_b^{-1}(y))v(y),$$

that is

$$a \succsim b \iff u(a,\sigma) := \sum_{s \in S} \sigma(s)u(a,s) \geq \sum_{s \in S} \sigma(s)u(b,s) =: u(b,\sigma).$$

4 Uncertainty: subjective expected utility

Suppose DM does not know the objective probabilities of states, i.e., that we have a problem of decision under uncertainty. According to the subjective expected utility theory, DM has quantifiable beliefs about such unknown probabilities. We can take a direct approach and just look at the subjective probabilities of states: if $\bar{\mu} \in \Delta(S)$ represents the subjective belief (or "conjecture") of DM, then he ranks mixed actions according to the subjective expected utility (SEU) criterion, i.e., for each $\alpha, \beta \in \Delta(A)$

$$\alpha \succsim \beta \iff \mathbb{E}_{\lambda_{\alpha,\bar{\mu}}}[v] \geq \mathbb{E}_{\lambda_{\beta,\bar{\mu}}}[v].$$

Specifically, recall that each mixed action $\alpha \in \Delta(A)$ determines a state-contingent lottery $\lambda_\alpha = (\lambda_{\alpha,s}) \in [\Delta(S)]^S$, that is, one lottery for each state s. Then

$$\begin{aligned}
\mathbb{E}_{\lambda_{\alpha,\bar{\mu}}}[v] &= \sum_{s \in S} \bar{\mu}(s)\mathbb{E}_{\lambda_{\alpha,s}}[v] = \sum_{s \in S} \bar{\mu}(s) \sum_{a \in A} \alpha(a)v(g(a,s)) \\
&= \sum_{s \in S}\sum_{a \in A} \bar{\mu}(s)\alpha(a)u(a,s).
\end{aligned}$$

Hence, for every $\alpha, \beta \in \Delta(A)$

$$\alpha \succsim \beta \iff \sum_{s \in S}\sum_{a \in A} \bar{\mu}(s)\alpha(a)u(a,s) \geq \sum_{s \in S}\sum_{a \in A} \bar{\mu}(s)\beta(a)u(a,s),$$

[5] For each (finite) space X, real-valued function $f \in \mathbb{R}^X$ and probability measure $\mu \in \Delta(X)$, we write

$$\mathbb{E}_\mu[f] := \sum_{x \in X} \mu(x)f(x).$$

or, using our abbreviations,

$$\alpha \succsim \beta \Longleftrightarrow u(\alpha, \bar{\mu}) \geq u(\beta, \bar{\mu}).$$

But it is more instructive to take an indirect approach. Let us use again the metaphor of the urn. If DM does not know the true composition σ of the urn, then he forms subjective beliefs about it. To avoid mathematical complications, we assume that DM deems possible a *finite* subset $\Sigma \subset \Delta(S)$ of urn compositions. Then his belief is represented by a subjective probability measure $\mu \in \Delta(\Sigma)$. For each urn composition $\sigma \in \Sigma$, a mixed action $\alpha \in \Delta(A)$ yields a lottery $\lambda_{\alpha,\sigma}$ and a corresponding objective expected utility

$$\mathbb{E}_{\lambda_{\alpha,\sigma}}[v] = u(\alpha, \sigma) = \sum_{s \in S} \sum_{a \in A} \sigma(s)\alpha(a)u(a, s).$$

DM ranks mixed actions according to the subjective expected value of his objective expected utility: for every $\alpha, \beta \in \Delta(A)$

$$\begin{aligned}
\alpha \quad \succsim \quad \beta &\Longleftrightarrow \sum_{\sigma \in \Sigma} \mu(\sigma)u(\alpha, \sigma) \geq \sum_{\sigma \in \Sigma} \mu(\sigma)u(\beta, \sigma) \\
&\Longleftrightarrow \sum_{\sigma \in \Sigma} \sum_{s \in S} \mu(\sigma)\sigma(s)u(\alpha, s) \geq \sum_{\sigma \in \Sigma} \sum_{s \in S} \mu(\sigma)\sigma(s)u(\beta, s) \\
&\Longleftrightarrow \sum_{s \in S} u(\alpha, s) \sum_{\sigma \in \Sigma} \mu(\sigma)\sigma(s) \geq \sum_{s \in S} u(\beta, s) \sum_{\sigma \in \Sigma} \mu(\sigma)\sigma(s),
\end{aligned}$$

where the second equivalence follows by replacing $u(\alpha, \sigma)$ [respectively $u(\beta, \sigma)$] with $\sum_{s \in S} \mu(\sigma)\sigma(s)u(\alpha, s)$ [respectively $\sum_{s \in S} \mu(\sigma)\sigma(s)u(\beta, s)$].

DM's belief μ on urn compositions induces subjective probabilities $\bar{\mu} \in \Delta(S)$ of the colors of the balls in the urn, called the **predictive probabilities** of states (colors), which are just the weighted average of the objective probabilities with weights given by belief μ:

$$\forall s \in S, \ \bar{\mu}(s) = \sum_{\sigma \in \Sigma} \mu(\sigma)\sigma(s).$$

With this, using the previous equivalence we obtain

$$\begin{aligned}
\alpha \quad \succsim \quad \beta &\Longleftrightarrow \sum_{s \in S} u(\alpha, s) \sum_{\sigma \in \Sigma} \mu(\sigma)\sigma(s) \geq \sum_{s \in S} u(\beta, \sigma) \sum_{\sigma \in \Sigma} \mu(\sigma)\sigma(s) \\
&\Longleftrightarrow \sum_{s \in S} \bar{\mu}(s)u(\alpha, s) \geq \sum_{s \in S} \bar{\mu}(s)u(\beta, s) \Longleftrightarrow u(\alpha, \bar{\mu}) \geq u(\beta, \bar{\mu}).
\end{aligned}$$

Hence we are back to the previous formulation.

Since DM only cares about the objective probabilities of consequences in each state, indeed his objective expected utility in each state, he actually ranks vectors $(u(\alpha, s))_{s \in S} \in \mathbb{R}^S$. If $(u(\alpha, s))_{s \in S} = (u(\beta, s))_{s \in S}$ then α and β are equivalent and represented by the same vector in $\mathbb{R}^S$. Such vectors are a kind of "reduced form"of Anscombe–Aumann acts.

The set of feasible vectors corresponding to mixed actions is the convex hull of the finite set of vectors corresponding to pure actions

$$\bar{A} = \text{convhull}\{(u(a, s))_{s \in S} : a \in A\},$$

a convex and compact polyhedron in $\mathbb{R}^S$ (see Chapter 3 of the lecture notes [3]). We will often refer to the vector in $\mathbb{R}^S$ corresponding to a mixed action $\alpha \in \Delta(A)$ under the surjection from $\Delta(A)$ to $\bar{A}$ defined by

$$\alpha \mapsto (u(\alpha, s))_{s \in S}.$$

Since (mixed) actions inducing the same utility vector are equivalent, we can represent DM's preferences within space $\mathbb{R}^S$ (which contains the feasible set $\bar{A}$) rather than space $\mathbb{R}^A$ (which contains the feasible set $\Delta(A)$).

Under the subjective expected utility theory, the DM's indifference sets in $\mathbb{R}^S$ are hyperplanes of the form

$$\left\{ \mathbf{U} \in \mathbb{R}^S : \sum_{s \in S} U(s)\bar{\mu}(s) = k \right\}$$

for some constant $k \in \mathbb{R}$, i.e., hyperplanes with normal vector $\bar{\mu}$. When $S = \{\ell, r\}$, $\bar{A}$ is a polygon and we have indifference lines in $\mathbb{R}^{\{\ell, r\}}$ defined by the equation

$$\bar{\mu}(\ell)U(\ell) + \bar{\mu}(r)U(r) = k,$$

i.e., lines with negative slope

$$\frac{dU(r)}{dU(\ell)} = -\frac{\bar{\mu}(\ell)}{\bar{\mu}(r)} = -\frac{\bar{\mu}(\ell)}{1 - \bar{\mu}(\ell)}.$$

A subjectively optimal mixed action α^* corresponds to either a vertex of $\bar{A}$, which implies that α^* is either a pure action or a convex combination of equivalent pure actions, or a point in the relative interior of an edge or face of the compact and convex polyhedron $\bar{A}$. Suppose that an optimal point $(u(\alpha^*, s))_{s \in S}$ is in the relative interior of a "full face" F (a set of dimension $|S| - 1$, e.g. a polygon when $|S| = 3$, or a line segment when $|S| = 2$); then face F must be a subset of the indifference set

$$H(\alpha^*, \bar{\mu}) := \left\{ \mathbf{U} \in \mathbb{R}^S : \sum_{s \in S} U(s)\bar{\mu}(s) = u(\alpha^*, \bar{\mu}) \right\},$$

which is a hyperplane; with this, every other vector in face F is also optimal. Indeed, face F is the convex hull of the vectors corresponding to optimal pure actions (see Lemma 1 in Chapter 3 of [3]):

$$F = \text{convhull}\{(u(a, s))_{s \in S} : u(a, \bar{\mu}) = u(\alpha^*, \bar{\mu})\}.$$

In the two-dimensional case with $S = \{\ell, r\}$, an optimal (pure or) mixed action corresponds to a point on the "North-East boundary" of the feasible set $\bar{A} \subseteq \mathbb{R}^{\{\ell, r\}}$. If this point is in the relative interior of an edge, then this edge must have slope $\bar{\mu}(\ell)/\bar{\mu}(r)$ in absolute value and its two endpoints (vertices) correspond to optimal pure actions. More generally, every optimal point $(u(\alpha^*, s))_{s \in S}$ must be on the part of the boundary of the polyhedron $\bar{A}$ with positive orientation, because $H(\alpha^*, \bar{\mu})$, the hyperplane through point $(u(\alpha^*, s))_{s \in S}$ with orientation $\bar{\mu}$, defines a half-space containing $\bar{A}$. Note that the boundary of $\bar{A}$ with positive orientation corresponds to the set of undominated mixed actions.

To sum up, *a SEU maximizer cannot have a strict preference for randomization*, for such a SEU-DM mixed actions are redundant.

5 Uncertainty aversion and preference for randomization

Consider a problem of decision under uncertainty. Let $\Sigma \subseteq \Delta(S)$ be the finite set of urn compositions deemed possible, with $|\Sigma| \geq 2$ For example, it may be known that there 10 balls in a two-color urn and the number of red balls is between 4 and 6:

$$\Sigma = \left\{\sigma \in \Delta(\{\ell, r\}) : \sigma(r) \in \left\{\frac{4}{10}, \frac{5}{10}, \frac{6}{10}\right\}\right\}.$$

Let us compare two types of DM. The first, DM', is subjectively certain that the true objective distribution is a particular $\hat{\sigma}$ (e.g., the uniform on S), i.e., he has a degenerate belief $\mu' = \delta_{\hat{\sigma}}$, where[6]

$$\delta_{\hat{\sigma}}(\sigma) = \begin{cases} 1, & \text{if } \sigma = \hat{\sigma}, \\ 0, & \text{if } \sigma \neq \hat{\sigma}. \end{cases}$$

The second one, DM", has a belief $\mu'' \in \Delta(\Sigma)$ with $|\mathrm{supp}\mu''| \geq 2$ that yields predictive probabilities of states $\bar{\mu}'' = \hat{\sigma}$. Since $|\mathrm{supp}\mu''| \geq 2$, DM" assigns positive subjective probability to at least two possible urn compositions in Σ, hence DM" is subjectively uncertain. However, in both cases the predictive probabilities are equal to $\widehat{\sigma}$: $\bar{\mu}' = \widehat{\sigma} = \bar{\mu}''$. According to the SEU maximization theory there is no difference between DM' and DM" because only the predictive probabilities matter for a SEU maximizer.

But there are other plausible decision models that predict different behavior for DM' and DM." According to such models, a decision maker – other things being equal – dislikes (mixed or pure) actions that give rise to an unknown objective expected utility, i.e., actions $\alpha \in \Delta(A)$ such that

$$u(\alpha, \sigma') \neq u(\alpha, \sigma'')$$

for some $\sigma', \sigma'' \in \Sigma$.

Consider the previous example, where Σ has three elements. Let $A = \{t, m, b\}$ (top, medium, bottom) and let the payoff function be given by the following matrix

$g(\cdot,\cdot)$	ℓ	r
t	0	3
m	1	1
b	3	0

Suppose that DM' is certain that there are exactly 5 red balls out of 10 in the urn: $\mu' = \delta_{\hat{\sigma}}$ with $\hat{\sigma}(r) = \frac{5}{10}$, whereas DM" deems the three urn compositions equally likely:

$$\mu''(\sigma(r) = 4/10) = \mu''(\sigma(r) = 5/10) = \mu''(\sigma(r) = 6/10).$$

Then $\bar{\mu}'(r) = \hat{\sigma}(r) = \frac{1}{2} = \bar{\mu}''(r)$. According to SEU theory DM' and DM" have the same set of optimal mixed actions, that is

$$\{\alpha \in \Delta(A) : \alpha(m) = 0\}.$$

[6] The Greek letter δ stands for "Dirac measure," but it also reminds the reader of "deterministic belief."

This set is isomorphic to the line segment in $\mathbb{R}^{\{\ell,r\}}$ with endpoints $(0,3)$ (top, t) and $(3,0)$ (bottom, b), i.e., the "North-East"edge of $\bar{A}$. In particular, it contains the mixed action α^* that assigns 50% chances to top and bottom, that is, the convex combination with equal weights of the two deterministic measures δ_t and δ_b:

$$\alpha^* = \frac{1}{2}\delta_t + \frac{1}{2}\delta_b.$$

As shown in Figure 1, mixed action α^* is the only one in the set $\{\alpha \in \Delta(A) : \alpha(m) = 0\}$ with the notable feature of yielding the *same objective expected utility (payoff) for each possible distribution*: $u(\alpha^*, \sigma) = \frac{3}{2}$ for each $\sigma \in \Sigma$.

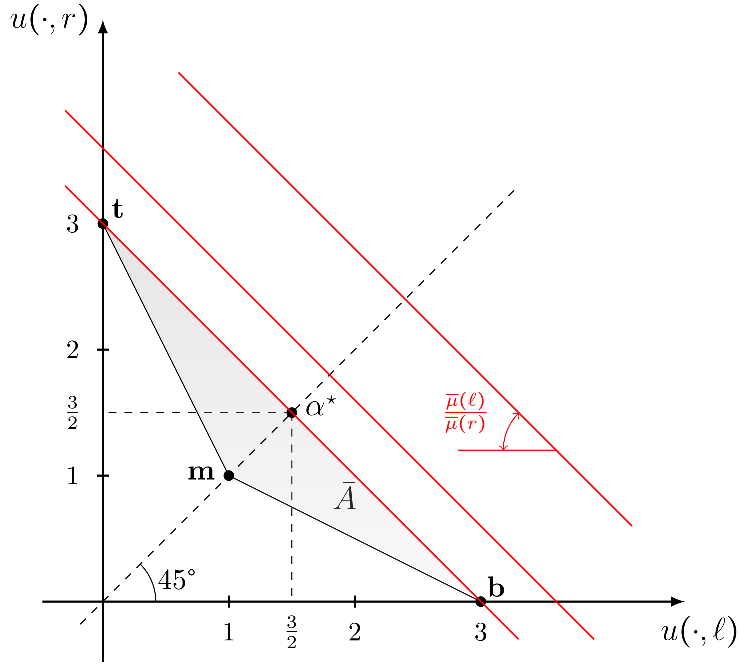

Figure 1: Optimal choice for SEU maximizer

If a decision maker dislikes choices that yield an uncertain objective expected payoff, this gives an edge to α^*. In particular, DM' who is subjectively certain that the true urn contains exactly 5 red balls, is still indifferent between all the mixed actions in $\{\alpha \in \Delta(A) : \alpha(m) = 0\}$. But DM", who is subjectively uncertain about the urn composition, should strictly prefer α^*. Thus, the models considered here may give a *strict preference for randomization.* The intuition is that randomization allows to "hedge against uncertainty" by shifting objective expected utility from one state to another so as to decrease the differences in expected utility across possible distributions.

We present two such models. In the first one, only the support $\operatorname{supp}\mu$ of the subjective belief $\mu \in \Delta(\Sigma)$ plays a role, whereas in the second one the whole belief μ plays a role. These models have the common feature that the decision maker, other things being equal, dislikes mixed actions that give rise to an uncertain objective expected payoff. Another way to say this is the following. Call mixed action $\alpha \in \Delta(A)$ **ambiguous** (relative to $\mathrm{supp}\mu$) if there are $\sigma', \sigma'' \in \mathrm{supp}\mu$ such that

$$u(\alpha, \sigma') \neq u(\alpha, \sigma'').$$

Note that an ambiguous action must yield consequences whose objective probabilities are uncertain. According to the decision criteria presented below, DM tends to dislike ambiguous actions. Therefore these decision criteria are also called **models of "ambiguity aversion"**.

5.1 The maxmin criterion

To ease notation, for each belief $\mu \in \Delta(\Sigma)$, let Σ_μ denote its support

$$\Sigma_\mu = \mathrm{supp}\mu = \{\sigma \in \Sigma : \mu(\sigma) > 0\}.$$

According to the **maxmin criterion** due to Gilboa and Schmeidler [8],[7] for every $\alpha, \beta \in \Delta(A)$, DM prefers α to β if and only if

$$\min_{\sigma \in \Sigma_\mu} u(\alpha, \sigma) \geq \min_{\sigma \in \Sigma_\mu} u(\beta, \sigma).$$

Intuitively, DM behaves *as if* he thought that he is playing against an adversarial Nature that, instead of choosing the state $s \in S$, chooses the probabilities of states σ within the restricted set Σ_μ *after* observing the DM's mixed action and with the goal of *minimizing* the expected payoff of DM. Obviously, this is just a metaphor for a criterion that captures the preference for those mixed actions that tend to equalize expected payoffs across distributions in Σ_μ, i.e., a preference for unambiguous actions.

Going back to our example of the two decision makers, for DM' we have the singleton $\Sigma_{\mu'} = \{\hat{\sigma}\}$ with $\hat{\sigma}(r) = \frac{1}{2}$, therefore

$$\begin{aligned} \arg \max_{\alpha \in \Delta(\{t,m,b\})} \min_{\sigma \in \Sigma_{\mu'}} u(\alpha, \sigma) &= \\ \arg \max_{\alpha \in \Delta(\{t,m,b\})} \left[\frac{1}{2} u(\alpha, \ell) + \frac{1}{2} u(\alpha, r)\right] &= \{\alpha \in \Delta(A) : \alpha(m) = 0\}. \end{aligned}$$

In other words, since DM' is certain that the true composition of the urn is 50%-50%, he behaves like a SEU maximizer with predictive probabilities $\bar{\mu}(\ell) = \frac{1}{2} = \bar{\mu}(r)$ (see Figure 1).

On the other hand, DM" has a belief μ'' that deems possible three distributions: $\Sigma_{\mu''} = \{\sigma^\ell, \hat{\sigma}, \sigma^r\}$ with $\sigma^\ell(r) = \frac{2}{5}$, $\hat{\sigma}(r) = \frac{1}{2}$, $\sigma^r(r) = \frac{3}{5}$. Hence each action α gives rise to three

[7] Gilboa and Schmeidler present axioms about preferences $\succeq$ over the set $[\Delta(Y)]^S$ of Anscombe–Aumann acts that yield such a functional-form representation: the DM satisfies the axioms if and only if there is a VNM utility function $v : Y \to \mathbb{R}$ and a (compact and convex) set of measures $\hat{\Sigma} \subseteq \Delta(S)$ such that

$$\alpha \succeq \beta \Leftrightarrow \min_{\sigma \in \hat{\Sigma}} \mathbb{E}_{\alpha,\sigma}[v] \geq \min_{\sigma \in \hat{\Sigma}} \mathbb{E}_{\beta,\sigma}[v]$$

Note that these authors do not explicitly mention subjective probabilistic beliefs μ, because such beliefs cannot be elicited from the preference relation. Here we interpret $\hat{\Sigma}$ as the convex hull of $\mathrm{supp}\mu$.

possible objective expected payoffs

$$\begin{aligned} u(\alpha,\sigma^{\ell}) &= \frac{3}{5}u(\alpha,\ell)+\frac{2}{5}u(\alpha,r), \\ u(\alpha,\hat{\sigma}) &= \frac{1}{2}u(\alpha,\ell)+\frac{1}{2}u(\alpha,r), \\ u(\alpha,\sigma^{r}) &= \frac{2}{5}u(\alpha,\ell)+\frac{3}{5}u(\alpha,r), \end{aligned}$$

and DM" solves

$$\max_{\alpha\in\Delta(\{t,m,b\})}\min\left\{\frac{3}{5}u(\alpha,\ell)+\frac{2}{5}u(\alpha,r),\frac{1}{2}u(\alpha,\ell)+\frac{1}{2}u(\alpha,r),\frac{2}{5}u(\alpha,\ell)+\frac{3}{5}u(\alpha,r)\right\}.$$

The indifference curves in $\mathbb{R}^{\{\ell,r\}}$ are piecewise linear with a kink on the 45°-line, slope $-\frac{3}{2}$ above it and slope $-\frac{2}{3}$ below it, where the ascissa is the expected payoff in state ℓ. Note that the same is true if the Decision Maker deems possible only the extreme distributions σ^{ℓ} and σ^{r}. The feasible set in $\mathbb{R}^{\{\ell,r\}}$ is

$$\bar{A}=\text{convhull}\{(0,3),(1,1),(3,0)\}.$$

The highest indifference curve in $\mathbb{R}^{\{\ell,r\}}$ with nonempty intersection with the feasible set $\bar{A}$ is the one with kink at point $(1.5,1.5)$, which corresponds to the mixed action $\alpha^{*}=\frac{1}{2}\delta_t+\frac{1}{2}\delta_b$. Thus, as shown in Figure 2, $\alpha^{*}\in\{\alpha\in\Delta(A):\alpha(m)=0\}$ is the only optimal action for DM", who has a strict preference for randomization.

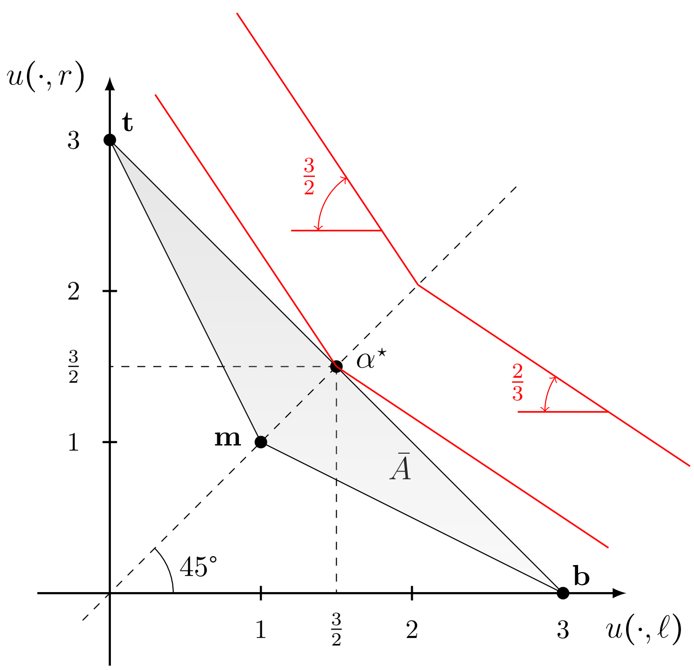

Figure 2: Optimal choice with the maxmim criterion

5.2 The smooth criterion

Now we consider a choice criterion due to Klibanoff, Marinacci and Mukerji [9] that makes use of the whole subjective belief $\mu \in \Delta(\Sigma)$ and yields smooth indifference sets in $\mathbb{R}^S$. Let

$$\underline{v} = \min_{\lambda \in \Delta(Y)} \mathbb{E}_\lambda[v],\ \bar{v} = \max_{\lambda \in \Delta(Y)} \mathbb{E}_\lambda[v].$$

DM's preferences are characterized by a continuous function

$$\phi : [\underline{v}, \bar{v}] \to \mathbb{R}$$

twice differentiable on the open interval $(\underline{v}, \bar{v})$ and such that $\phi' > 0$ (ϕ is strictly increasing) and $\phi'' \le 0$ (ϕ is concave). This function ϕ is a kind of **second-order utility function**: for each objective expected payoff $u(\alpha, \sigma)$ (a VNM utility) we may interpret $\phi(u(\alpha, \sigma))$ as the "second-order utility" of $u(\alpha, \sigma)$. According to the following criterion a DM with belief $\mu \in \Delta(\Sigma)$ and characterized by ϕ wants to maximize the subjective expected value of his "second-order utility":

$$\alpha \succsim \beta \Leftrightarrow \mathbb{E}_\mu[\phi(u(\alpha, \sigma))] \ge \mathbb{E}_\mu[\phi(u(\beta, \sigma))],$$

that is,

$$\alpha \succsim \beta \Leftrightarrow \sum_{\sigma \in \Sigma} \mu(\sigma)\phi(u(\alpha, \sigma)) \ge \sum_{\sigma \in \Sigma} \mu(\sigma)\phi(u(\beta, \sigma)).$$

Note that when $\phi'' = 0$ (i.e., ϕ is affine) we obtain the standard SEU criterion: in this case $\phi(U) = mU + n$ with $m = \phi'(U) > 0$. Therefore, for every $\alpha \in \Delta(A)$,

$$\sum_{\sigma \in \Sigma} \mu(\sigma)\phi(u(\alpha, \sigma)) = \sum_{\sigma \in \Sigma} \mu(\sigma)(mu(\alpha, \sigma) + n) = m \sum_{\sigma \in \Sigma} \mu(\sigma)u(\alpha, \sigma) + n$$

and

$$\alpha \succsim \beta \Leftrightarrow \sum_{\sigma \in \Sigma} \mu(\sigma)u(\alpha, \sigma) \ge \sum_{\sigma \in \Sigma} \mu(\sigma)u(\beta, \sigma).$$

As we have seen, for a SEU maximizer only the predictive probabilities of states $\bar{\mu} \in \Delta(S)$ matter [where $\bar{\mu}(s) = \sum_{\sigma \in \Sigma} \mu(\sigma)\sigma(s)$] and the indifference sets in $\mathbb{R}^S$ are hyperplanes with normal vector $\bar{\mu}$. But if $\phi'' < 0$ then the whole belief $\mu \in \Delta(\Sigma)$ matters and the indifference sets in $\mathbb{R}^S$ are smooth surfaces with orientation $\bar{\mu}$ only on the main diagonal and strictly convex upper contour sets.

To ease notation, let us write $\sigma_s = \sigma(s)$ for the objective probability of s under distribution $\sigma \in \Sigma$, $\bar{\mu}_s = \bar{\mu}(s)$ for the predictive probability of s, and $\mathbf{U} = (U_s)_{s \in S} \in \mathbb{R}^S$ for a generic vector of state-contingent expected payoffs. Define,

$$V(\mathbf{U}, \mu) = \sum_{\sigma \in \Sigma} \mu(\sigma)\phi\left(\sum_{s \in S} \sigma_s U_s\right)$$

as the value of state-contingent expected payoff $\mathbf{U}$ given belief μ. Along an indifference surface in $\mathbb{R}^S$ the value $V(\mathbf{U}, \mu)$ is constant, hence the differential dV is zero:

$$dV = \sum_{s \in S} \frac{\partial V}{\partial U_s} dU_s = 0.$$

The orientation of an indifference surface is given by the gradient of V:

$$\nabla V = \left(\frac{\partial V}{\partial U_s}\right)_{s\in S} = \left(\sum_{\sigma\in\Sigma}\phi'\left(\sum_{s\in S}\sigma_s U_s\right)\mu(\sigma)\sigma_s\right)_{s\in S}.$$

On the main diagonal, i.e., if $U_s = U$ for every s and some $U \in \mathbb{R}$, the gradient of V is proportional to the vector $\bar{\mu}$ of predictive probabilities:

$$\nabla V = \left(\phi'(U)\sum_{\sigma\in\Sigma}\mu(\sigma)\sigma_s\right)_{s\in S} = \phi'(U)\left(\bar{\mu}_s\right)_{s\in S},$$

where we used the fact that $\sum_{s\in S}\sigma_s = 1$, hence $\sum_{s\in S}\sigma_s U_s = U$, and $\sum_{\sigma\in\Sigma}\mu(\sigma)\sigma_s = \bar{\mu}_s$.

Now consider the two-states case, $S = \{\ell, r\}$. By the implicit function theorem,

$$\left.\frac{dU_r}{dU_\ell}\right|_{V=\text{const.}} = -\frac{\frac{\partial V}{\partial U_\ell}}{\frac{\partial V}{\partial U_r}} = -\frac{\sum_{\sigma\in\Sigma}\phi'(\sigma_\ell U_\ell + \sigma_r U_r)\mu(\sigma)\sigma_\ell}{\sum_{\sigma\in\Sigma}\phi'(\sigma_\ell U_\ell + \sigma_r U_r)\mu(\sigma)\sigma_r} < 0$$

where the denominator is positive if $\mu(\sigma)\sigma_r > 0$ for some $\sigma \in \Sigma$, i.e., if $\sum_{\sigma\in\Sigma}\mu(\sigma)\sigma_r = \bar{\mu}_r > 0$ (because $\phi' > 0$). On the 45°-line $U_\ell = U_r = U$ and

$$\left.\frac{dU_r}{dU_\ell}\right|_{V=\text{const.}} = -\frac{\sum_{\sigma\in\Sigma}\phi'(U)\mu(\sigma)\sigma_\ell}{\sum_{\sigma\in\Sigma}\phi'(U)\mu(\sigma)\sigma_r} = -\frac{\sum_{\sigma\in\Sigma}\mu(\sigma)\sigma_\ell}{\sum_{\sigma\in\Sigma}\mu(\sigma)\sigma_r} = -\frac{\bar{\mu}_\ell}{\bar{\mu}_r}$$

(see Figure 3).

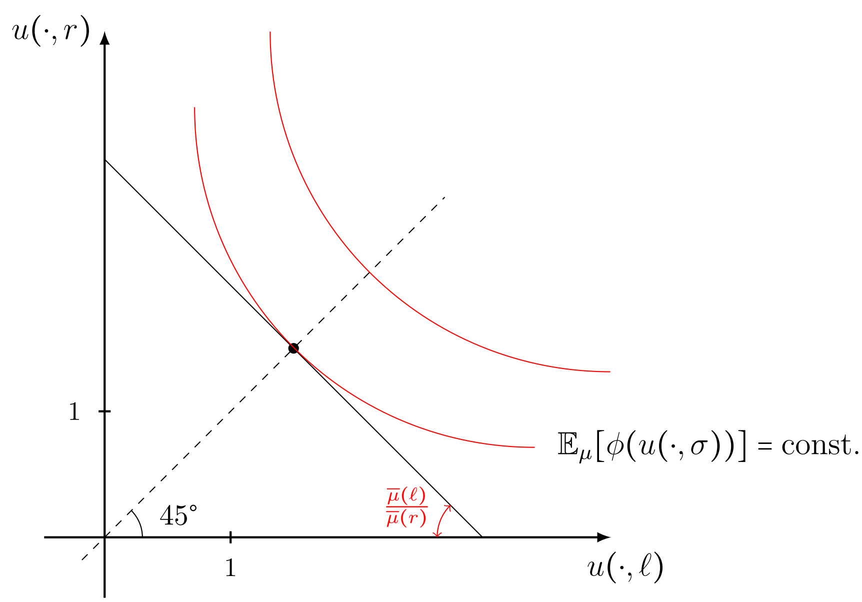

Figure 3: Indifference curves with the smooth criterion

Intuitively, when the expected payoff is independent of the state, then it is also independent of the objective distribution σ; therefore the ambiguity aversion caused by uncertainty about σ and captured by ϕ does not matter. Hence, in a neighborhood of a point on the 45°-line, the value of an expected payoff vector $\mathbf{U} = (U_s)_{s\in S}$ is well approximated by the subjective expected value

$$\sum_{\sigma\in\Sigma}\mu(\sigma)\sum_{s\in S}\sigma_s U_s = \sum_{s\in S}U_s\sum_{\sigma\in\Sigma}\mu(\sigma)\sigma_s = \sum_{s\in S}\bar{\mu}_s U_s.$$

Finally, we show that, in the two-states case, strict concavity of ϕ implies that the upper contour sets $\{\mathbf{U} : V(\mathbf{U},\mu) \geq \bar{V}\}$ are *strictly* convex as long as DM is uncertain, i.e., if he assigns strictly positive probability to at least two models ($|\mathrm{supp}\mu| \geq 2$).

First we show that *concavity of ϕ implies concavity of $V(\mathbf{U},\mu)$ in* $\mathbf{U}$ for each μ. Then we derive strict convexity of the upper contour sets in the two-states case.

Take two vectors $\mathbf{U}', \mathbf{U}'' \in \mathbb{R}^S$ and a weight $\gamma \in [0,1]$. Then

$$\begin{aligned}
V(\gamma\mathbf{U}' + (1-\gamma)\mathbf{U}'',\mu) &= \sum_{\sigma\in\Sigma}\mu(\sigma)\phi\left(\sum_{s\in S}\sigma_s(\gamma U'_s + (1-\gamma)U''_s)\right)\\
&= \sum_{\sigma\in\Sigma}\mu(\sigma)\phi\left(\gamma\sum_{s\in S}\sigma_s U'_s + (1-\gamma)\sum_{s\in S}\sigma_s U''_s\right)\\
&\geq \sum_{\sigma\in\Sigma}\mu(\sigma)\left(\gamma\phi\left(\sum_{s\in S}\sigma_s U'_s\right) + (1-\gamma)\phi\left(\sum_{s\in S}\sigma_s U''_s\right)\right)\\
&= \gamma\sum_{\sigma\in\Sigma}\mu(\sigma)\left(\phi\left(\sum_{s\in S}\sigma_s U'_s\right)\right) + (1-\gamma)\sum_{\sigma\in\Sigma}\mu(\sigma)\left(\phi\left(\sum_{s\in S}\sigma_s U''_s\right)\right)\\
&= \gamma V(\mathbf{U}',\mu) + (1-\gamma)V(\mathbf{U}'',\mu),
\end{aligned}$$

where the inequality follows from the concavity of ϕ.

Now consider the two-state case and assume that $\mathbf{U}' \neq \mathbf{U}''$ and $0 < \gamma < 1$. We show below that $\sum_{s\in S}\sigma_s U'_s = \sum_{s\in S}\sigma_s U''_s$ for *at most* one σ. Since $\mathrm{supp}\mu$ contains at least two elements, it follows that there must be at least one $\hat{\sigma} \in \mathrm{supp}\mu$ such that

$$\sum_{s\in S}\hat{\sigma}_s U'_s \neq \sum_{s\in S}\hat{\sigma}_s U''_s.$$

Since ϕ is strictly concave and $0 < \gamma < 1$, then

$$\phi\left(\gamma\sum_{s\in S}\hat{\sigma}_s U'_s + (1-\gamma)\sum_{s\in S}\hat{\sigma}_s U''_s\right) > \gamma\phi\left(\sum_{s\in S}\hat{\sigma}_s U'_s\right) + (1-\gamma)\phi\left(\sum_{s\in S}\hat{\sigma}_s U''_s\right).$$

Hence, the weak inequality that we proved for the general case must be strict, showing that

$$V(\gamma\mathbf{U}' + (1-\gamma)\mathbf{U}'',\mu) > \max\{V(\mathbf{U}',\mu), V(\mathbf{U}'',\mu)\}.$$

Thus, the upper contour sets are strictly convex.

Now we give the proof that we postponed, i.e., we prove that there is at most one distribution $\sigma \in \Delta(\{\ell,r\})$ equalizing $\sum_{s\in S}\sigma_s U'_s$ and $\sum_{s\in S}\sigma_s U''_s$. Consider the unique straight line through $\mathbf{U}'$ and $\mathbf{U}''$. If it has positive slope, one of the two points dominates the other and there is no $\sigma \in \Delta(\{\ell,r\})$ that equalizes $\sum_{s\in S}\sigma_s U'_s$ and $\sum_{s\in S}\sigma_s U''_s$. Otherwise, the negative slope identifies the probability ratio σ_ℓ/σ_r, hence the unique $\sigma \in \Delta(\{\ell, r\})$ equalizing $\sum_{s\in S}\sigma_s U'_s$ and $\sum_{s\in S}\sigma_s U''_s$.

More formally, suppose that σ is such that

$$(1-\sigma_r)U'_\ell + \sigma_r U'_r = (1-\sigma_r)U''_\ell + \sigma_r U''_r$$

and assume w.l.o.g. that r has strictly positive probability, i.e., $\sigma_r > 0$. Then

$$\frac{1-\sigma_r}{\sigma_r} = \frac{U''_r - U'_r}{U'_\ell - U'_r}$$

$$\sigma_r = \frac{U'_\ell - U'_r}{U'_\ell - U'_r + U''_r - U'_r}.$$

Since $\Sigma \subseteq \Delta(\{\ell, r\})$, there is at most one element σ of Σ such that $\sum_{s\in S}\sigma_s U'_s = \sum_{s\in S}\sigma_s U''_s$.

To sum up, *in the two-state case, if ϕ is strictly concave and DM is uncertain, then the upper contour sets are strictly convex and there is always a unique solution to the problem*

$$\max_{\mathbf{U}\in\bar{A}} V(\mathbf{U}, \mu).$$

If this solution is not a vertex of the feasible set $\bar{A}$ and if the vectors $(u(a,s))_{s\in S}$ $(a \in A)$ are linearly independent, then DM strictly prefers to randomize (if the vectors were not linearly independent a solution that is not a vertex might be induced by a pure action).

Our leading example illustrates this. The "North-Est" boundary of the feasible set

$$\bar{A} = \text{convhull}\{(0,3),(1,1),(3,0)\}$$

is the segment that joins $(0,3)$ (bottom) to $(3,0)$ (top). Its slope is -1. The unique solution to problem $\max_{\mathbf{U}\in\bar{A}} V(\mathbf{U}, \mu)$ is the point $\mathbf{U}^*$ on this segment where the marginal rate of substitution between objective expected utility in state r and in state ℓ is one:

$$MRS_{r,\ell}(\mathbf{U}^*, \mu) := \frac{\frac{\partial V}{\partial U_\ell}(\mathbf{U}^*, \mu)}{\frac{\partial V}{\partial U_r}(\mathbf{U}^*, \mu)} = 1.$$

We know that on the diagonal the marginal rate of substitution is the ratio of the predictive probabilities:

$$U_\ell = U_r \Rightarrow MRS_{r,\ell}(\mathbf{U}, \mu) = \frac{\bar{\mu}(\ell)}{\bar{\mu}(r)}.$$

Consider the belief of DM' and the belief of DM". The belief of DM', μ', assigns probability 1 to the uniform distribution $\hat{\sigma} = \frac{1}{2}\delta_\ell + \frac{1}{2}\delta_r$, thus the predictive is $\bar{\mu} = \hat{\sigma}$, the indifference sets are parallel straight lines with slope -1, DM' behaves like a SEU maximizer and is indifferent between all the mixtures of top and bottom.

Now consider the subjectively uncertain decision maker DM". Although he has the same predictive $\bar{\mu}'' = \hat{\sigma}$, as argued above, his upper contour sets are strictly convex. Since $\bar{\mu}''(\ell) = \bar{\mu}''(r)$, the first-order condition $MRS_{r,\ell}(\mathbf{U}^*, \mu) = 1$ is satisfied only on the diagonal, at point $\mathbf{U}^* = (1.5, 1.5)$, which corresponds to the mixed action $\alpha^* = \frac{1}{2}\delta_t + \frac{1}{2}\delta_b$ (see Figure 4).

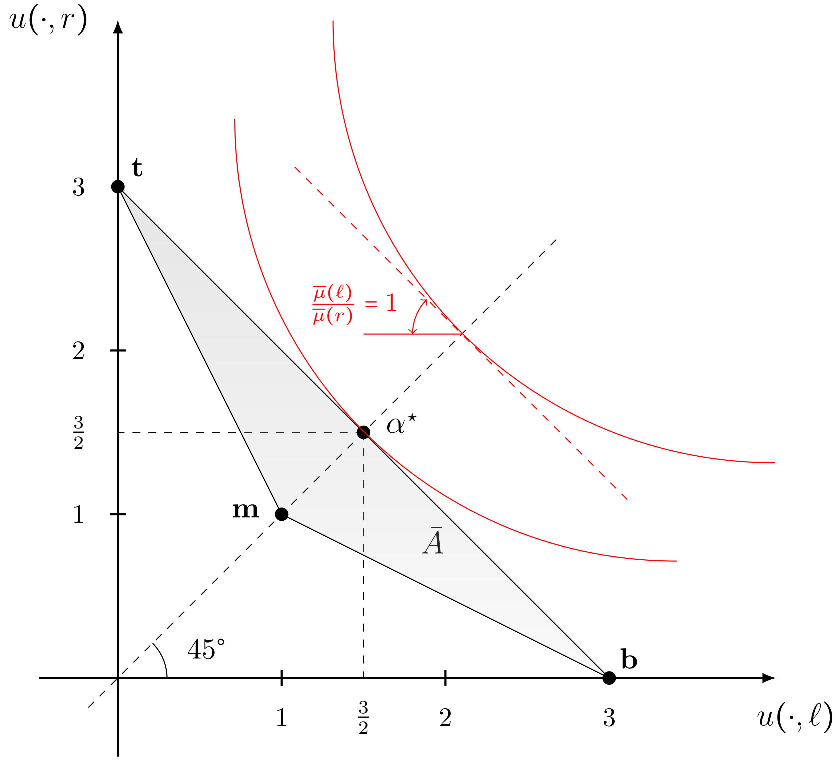

Figure 4: Optimal choice with the smooth criterion

6 Uncertainty aversion and selfconfirming equilibrium

In a **selfconfirming equilibrium** (SCE) each agent best responds to subjective beliefs confirmed by his personal experience, which in turn depends on his choice. Confirmed beliefs may be incorrect, therefore in an SCE agents' choices may be objectively suboptimal, even if they are subjectively optimizing; see [3, Chapter 5]. In the traditional notion of SCE, agents are SEU maximizers; but the SCE idea has been extended to allow for uncertainty aversion (see [4] and [5]).

Here we limit our attention to a DM in a recurrent decision problem who does not know the true probabilities of states. In a selfconfirming equilibrium his choice is subjectively optimal given his beliefs and the observed frequencies of consequences given that this choice is repeated over time match his subjective beliefs. This analysis can be used as a building block for a game-theoretic definition of SCE.

Let $\sigma^\circ \in \Delta(S)$ denote the true **objective distribution** (urn composition). In the recurrent decision problem, the state of nature is determined in each period with an i.i.d. draw with distribution σ°. DM knows that the draws are i.i.d., but does not know σ°. We consider an uncertainty averse DM who chooses according to the smooth criterion and is characterized by a concave second-order utility function ϕ. After each choice, DM obtains some **information feedback**, described by a function

$$f : A \times S \to M.$$

Here we assume for simplicity that *DM observes ex post the realized consequence* $y = g(a, s)$.

If there are *monetary consequences* ($Y \subseteq \mathbb{R}$) and the utility of wealth is strictly increasing, then we can assume w.l.o.g. that *the feedback function coincides with the payoff function*, i.e., $M = \mathbb{R}$ and $f = u$. We also assume, again for simplicity, that *if DM delegates his choice to a randomization device, he does not observe the realized action* a, he only observes the payoff $u(a,s)$ determined by the realized action a and by the realized state s.

Suppose that DM repeats the same (randomized) choice α infinitely many times. By the law of large numbers, the frequencies of monetary payoffs that he observes in the long run coincide with the probabilities of the lottery $\lambda_{\alpha,\sigma^\circ}$ induced by α and σ°. Hence he learns that the objective expected payoff of α is

$$u(\alpha,\sigma^\circ) = \sum_{a\in A}\sum_{s\in S} u(a,s).$$

But this does not mean that DM can identify the probabilities of states σ°. He finds out that the true probability distribution belong to the **partial identification set**

$$\begin{aligned}\hat{\Sigma}(\alpha,\sigma^\circ) \quad &: \quad = \{\sigma \in \Delta(S) : \lambda_{\alpha,\sigma} = \lambda_{\alpha,\sigma^\circ}\} \\ &= \quad \{\sigma \in \Delta(S) : \hat{u}(\alpha\times\sigma) = \hat{u}(\alpha\times\sigma^\circ)\}.\end{aligned}$$

[Recall that, for each function $f : X \to Y$, $\hat{f} : \Delta(X) \to \Delta(Y)$ is the corresponding push-forward function; here $\hat{u} : \Delta(A\times S) \to \Delta(Y)$ is the pushforward of $u : A\times S \to Y$, and $\alpha\times\sigma \in \Delta(A\times S)$ is the product measure obtained from the marginal measures $\alpha \in \Delta(A)$ and $\sigma \in \Delta(S)$.] For example, if the expected payoff $u(\alpha,s)$ is independent of state s, i.e., if there is some constant k such that $u(\alpha,s) = k$ for each $s \in S$, then DM remains in the dark because in this case $\hat{\Sigma}(\alpha,\sigma^\circ) = \Delta(S)$.

Note that, by definition, *expected payoff* – as a function of σ – *is constant* and equal to $u(\alpha,\sigma^\circ)$ *on the partially identified set* $\hat{\Sigma}(\alpha,\sigma^\circ)$. Intuitively, a choice-belief pair (α^*,μ) is an SCE if (1) α^* is subjectively optimal given μ and (2) μ is consistent with that fact that DM has "learned" the objective distribution of payoffs.

To give a more formal definition, we must be precise about the feasible set. In the previous sections we assumed that DM has a very rich commitment technology allowing him to span the whole set of mixed actions $\Delta(A)$. Here we generalize the analysis and assume that the *set of feasible mixed actions is some* compact *set* $\mathcal{A}$, with $A \subseteq \mathcal{A} \subseteq \Delta(A)$ (note, we identify A with the set of Dirac measures $\{\delta_a : a \in A\}$), and we let $\overline{\mathcal{A}} \subseteq \mathbb{R}^S$ denote the corresponding set of feasible (expected) payoff vectors. The two extreme cases are $\mathcal{A} = A$ and $\mathcal{A} = \Delta(A)$, both sets are compact, the second is also convex. With this, we say that (α^*,μ) is a **selfconfirming equilibrium** (given $\mathcal{A}$) if

1. (subjective rationality) $\alpha^* \in \arg\max_{\alpha\in\mathcal{A}} \mathbb{E}_\mu[\phi(u(\alpha,\sigma))]$,
2. (confirmed beliefs) $\mathrm{supp}\mu \subseteq \hat{\Sigma}(\alpha^*,\sigma^\circ)$.

Since $u(\alpha^*,\sigma) = u(\alpha^*,\sigma^\circ)$ for every $\sigma \in \hat{\Sigma}(\alpha^*,\sigma^\circ)$, the confirmed-beliefs condition implies that

$$\mu\left(\{\sigma : u(\alpha^*,\sigma) = u(\alpha^*,\sigma^\circ)\}\right) = 1.$$

Indeed, the confirmed-beliefs condition reflects the informal assumption that DM has "learned" the probabilities of monetary payoffs induced by choice α^*, and therefore he has also "learned" the objective expected payoff $u(\alpha^*, \sigma^\circ)$.

Is this notion of SCE different from the traditional one that assumes SEU maximization (or $\phi'' = 0$)? This depends on the feasible set $\mathcal{A}$. Indeed, one can use a separation argument to show that the following result: If $\mathcal{A}$ is convex,[8] the set of SCEs does not depend on the degree of ambiguity aversion. More formally, *for every pair* (α^*, μ), *if* (α^*, μ) *is an SCE for the concave second-order utility* ϕ, *then* (α^*, μ) *is an SCE for* every *concave* $\bar{\phi}$, *including every affine* $\bar{\phi}$.

To gain some intuition for this irrelevance result, suppose that

$$\mathcal{A} = \{\alpha \in \mathbb{R}^A : F_k(\alpha) \leq 0, k = 1, \dots K\}$$

where each F_k is a convex function differentiable in a neighborhood of the SCE choice α^*.[9] Then subjective optimality at α^* implies that the following first-order conditions hold:

$$\forall a \in A,\ \frac{\partial}{\partial \alpha_a}\left(\mathbb{E}_\mu[\phi(u(\alpha,\sigma))] - \sum_{k=1}^{K} \lambda_k F_k(\alpha)\right)\Bigg|_{\alpha=\alpha^*} = 0$$

where λ_k denotes the Lagrangian multiplier of the kth constraint. Now note that

$$\frac{\partial}{\partial \alpha_a} u(\alpha,\sigma) = \frac{\partial}{\partial \alpha_a}\left(\sum_{a' \in A} \alpha_{a'} u(a', \sigma)\right) = u(a,\sigma),$$

therefore

$$\frac{\partial}{\partial \alpha_a}\left(\mathbb{E}_\mu[\phi(u(\alpha,\sigma))]\right)\Big|_{\alpha=\alpha^*} = \sum_{\sigma \in \mathrm{supp}\mu} \mu(\sigma)\phi'(u(\alpha^*,\sigma))u(a,\sigma).$$

Since the belief-confirmation condition implies that $u(\alpha^*, \sigma) = u(\alpha^*, \sigma^\circ)$ for every $\sigma \in \mathrm{supp}\mu \subseteq \hat{\Sigma}(\alpha^*, \sigma^\circ)$, we get

$$\frac{\partial}{\partial \alpha_a}\left(\mathbb{E}_\mu[\phi(u(\alpha,\sigma))]\right)\Big|_{\alpha=\alpha^*} = \phi'(u(\alpha^*,\sigma^\circ)) \sum_{\sigma \in \mathrm{supp}\mu} \mu(\sigma)u(a,\sigma).$$

We can further simplify the expression recalling that

$$\begin{aligned}
\sum_{\sigma \in \mathrm{supp}\mu} \mu(\sigma)u(a,\sigma) &= \sum_{\sigma \in \mathrm{supp}\mu}\sum_{s \in S} \mu(\sigma)\sigma(s)u(a,s) \\
&= \sum_{s \in S} u(a,s) \sum_{\sigma \in \mathrm{supp}\mu} \mu(\sigma)\sigma(s) \\
&= \sum_{s \in S} u(a,s)\bar{\mu}(s) = u(a,\bar{\mu}).
\end{aligned}$$

Therefore

$$\frac{\partial}{\partial \alpha_a}\left(\mathbb{E}_\mu[\phi(u(\alpha,\sigma))]\right)\Big|_{\alpha=\alpha^*} = \phi'(u(\alpha^*,\sigma^\circ))u(a,\bar{\mu}).$$

[8] This sufficient condition is not tight. There are more general sufficient conditions that do not require the convexity of $\mathcal{A}$.

[9] These constraints either include the normalization and weak positivity constraints defining $\Delta(A)$, or make some or all of them redundant; whatever the case $\mathcal{A} \subseteq \Delta(A)$.

Hence, we obtain marginal rates of substitution (MRS) that do not depend on ϕ at (α^*, μ): for all $a, b \in A$

$$MRS_{a,b}(\alpha^*, \mu) := \frac{\frac{\partial}{\partial \alpha_a} \left(\mathbb{E}_\mu[\phi(u(\alpha, \sigma))]\right)\big|_{\alpha=\alpha^*}}{\frac{\partial}{\partial \alpha_b} \left(\mathbb{E}_\mu[\phi(u(\alpha, \sigma))]\right)\big|_{\alpha=\alpha^*}} = \frac{u(a, \bar{\mu})}{u(b, \bar{\mu})}.$$

The first-order conditions yield, for all $a, b \in A$,

$$MRS_{a,b}(\alpha^*, \mu) = \frac{u(a, \bar{\mu})}{u(b, \bar{\mu})} = \frac{\frac{\partial}{\partial \alpha_a} \left(\sum_{k=1}^K \lambda_k F_k(\alpha)\right)\big|_{\alpha=\alpha^*}}{\frac{\partial}{\partial \alpha_b} \left(\sum_{k=1}^K \lambda_k F_k(\alpha)\right)\big|_{\alpha=\alpha^*}} =: MRT_{a,b}(\alpha^*).$$

These MRS=MRT (marginal rate of transformation) conditions are independent of ϕ.

Now we illustrate the SCE analysis with our leading two-states example, where $S = \{\ell, r\}$ and u is given by (5). Suppose that $\sigma^\circ(r) = \frac{3}{5}$. Thus, the objectively optimal choice is t, top.

First note that if DM plays long enough a mixed action α^* such that $u(\alpha^*, \ell) \neq u(\alpha^*, r)$, the long-run frequencies of observed payoffs identify σ°. Indeed, the DM observes in the long run that his expected payoff is $u(\alpha, \sigma) = u(\alpha^*, \sigma^\circ)$. Thus

$$(1 - \sigma(r))u(\alpha^*, \ell) + \sigma(r)u(\alpha^*, r) = u(\alpha^*, \sigma^\circ).$$

Since $u(\alpha^*, r) - u(\alpha^*, \ell) \neq 0$, this yields

$$\sigma(r) = \frac{u(\alpha^*, \sigma^\circ) - u(\alpha^*, \ell)}{u(\alpha^*, r) - u(\alpha^*, \ell)} = \frac{\frac{3}{5}u(\alpha^*, r) + \frac{2}{5}u(\alpha^*, \ell) - u(\alpha^*, \ell)}{u(\alpha^*, r) - u(\alpha^*, \ell)} = \frac{3}{5}.$$

Therefore, the only SCE (α^*, μ) with $u(\alpha^*, \ell) \neq u(\alpha^*, r)$ (off the main diagonal) is the "Nash equilibrium" $(\alpha^*, \mu) = (\delta_t, \delta_{\sigma^\circ})$.[10]

If instead $u(\alpha^*, \ell) = u(\alpha^*, r)$, whether α^* is a SCE choice or not depends on the "fundamentals," $\mathcal{A}$ ("technology") and ϕ ("tastes"). Suppose, for example, that $\mathcal{A} = A$ (no randomization is feasible); then m is an SCE action for some belief μ provided that ϕ is sufficiently concave (see Figure 5). If instead ϕ is affine ($\phi'' = 0$), or close to an affine function, DM chooses as a SEU maximizer, hence he does not choose m, because m is strictly dominated by the unfeasible mixed action $\frac{1}{2}\delta_t + \frac{1}{2}\delta_b$ and hence it is not an SEU-best reply in A against any predictive belief $\bar{\mu}$.

At the other extreme, if $\mathcal{A} = \Delta(A)$ the result of irrelevance of the attitudes toward uncertainty (of function ϕ) applies. Indeed, either the equilibrium choice is revealing (off the diagonal), and we can only have the Nash equilibrium $(\alpha^*, \mu) = (\delta_t, \delta_{\sigma^\circ})$, or it is non-revealing (on the diagonal). In the second case, the optimality condition implies that α^* corresponds to the payoff vector at the intersection between the North-East boundary of $\overline{\mathcal{A}}$ and the diagonal, that is, $\alpha_t^* = \alpha_b^* = \frac{1}{2}$, $u(\alpha^*, \ell) = u(\alpha^*, r) = \frac{3}{2}$. The second equality means that the confirmation condition is satisfied: DM only observes in the long run that he "wins" (he gets 3 utils) half of the times, which is true under such mixed action α^* for every distribution σ. The subjective belief μ that supports such α^* as an optimal choice

[10] In a Nash equilibrium each agent best responds to correct beliefs. Hence our terminology makes sense.

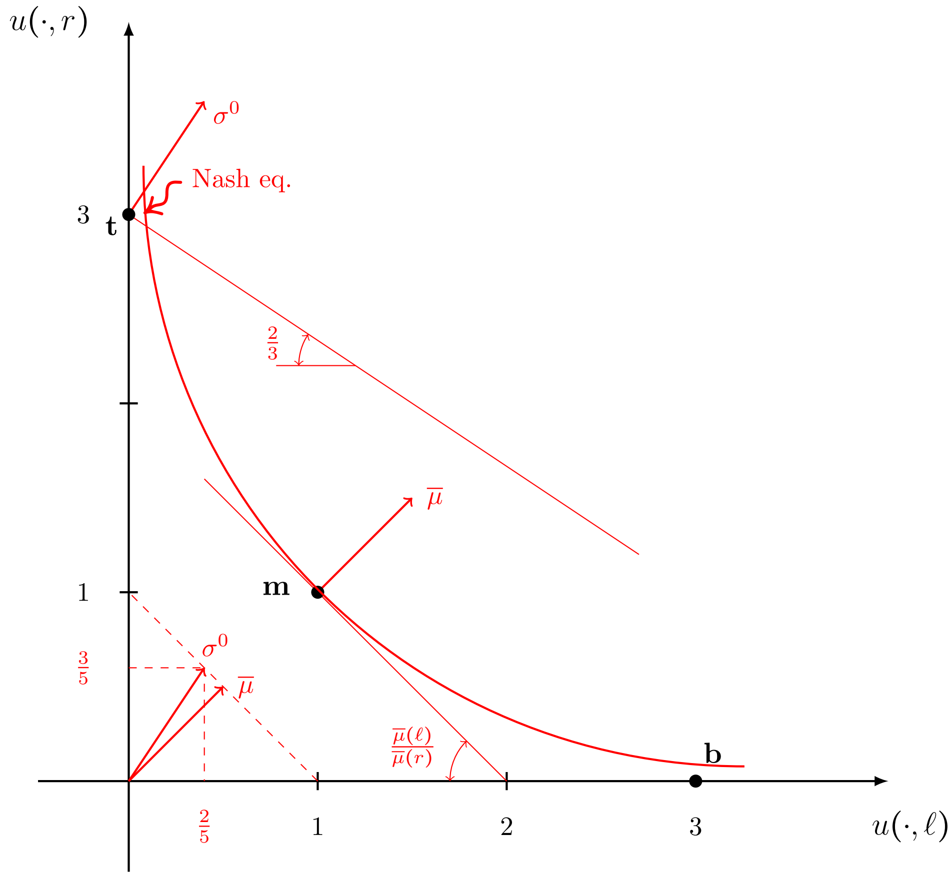

Figure 5: If ϕ is sufficiently concave and $\mathcal{A} = A$ then m is a SCE choice

(subjective rationality condition) must yield the uniform predictive distribution $\bar{\mu}(\ell) = \bar{\mu}(r)$. To see this, note that the tangency condition at the subjectively optimal choice is just

$$\frac{\bar{\mu}(\ell)}{\bar{\mu}(r)} = 1,$$

because for each point on the diagonal the slope of the indifference curve is $-\frac{\bar{\mu}(\ell)}{\bar{\mu}(r)}$, which is independent of ϕ (see Figure 3). Therefore, if $\mathcal{A} = \Delta(A)$, every (α^*, μ) such that $\alpha_t^* = \alpha_b^* = \bar{\mu}(r) = \frac{1}{2}$ is an SCE for every concave second-order utility ϕ.

References

[1] Anscombe, F.J. and R. Aumann (1963): "A definition of subjective probability," *Annals of Mathematical Statistics*, 34, 199-205.

[2] Battigalli, P. (2020): *Mathematical Language and Game Theory.* Typescript, Bocconi University.

[3] Battigalli, P., E. Catonini, and N. De Vito (2021): *Game Theory: Analysis of Strategic Thinking.* Typescript, Bocconi University.

[4] Battigalli P., S. Cerreia-Vioglio, F. Maccheroni, and M. Marinacci (2015): "Self-Confirming Equilibrium and Model Uncertainty," *American Economic Review*, 105, 646-677.

[5] Battigalli, P., S. Cerreia-Vioglio, F. Maccheroni, and M. Marinacci (2016): "Analysis of Information Feedback and Selfconfirming Equilibrium," *Journal of Mathematical Economics*, 66, 40–51.

[6] Battigalli, P., S. Cerreia-Vioglio, F. Maccheroni, and M. Marinacci (2017): "Mixed Extensions of Decision Problems Under Uncertainty," *Economic Theory,* 63, 827-866.

[7] Gilboa, I. and M. Marinacci (2013): "Ambiguity and the Bayesian paradigm," in D. Acemoglu, M. Arellano, and E. Dekel (eds.) *Advances in Economics and Econometrics: Theory and Applications,* Vol. 1, 179-242. Cambridge: Cambridge University Press.

[8] Gilboa, I. and D. Schmeidler (1989): "Maxmin expected utility with a non-unique prior," *Journal of Mathematical Economics*, 18, 141-153.

[9] Klibanoff P., M. Marinacci, and S. Mukerji (2005): "A smooth model of decision making under ambiguity," *Econometrica*, 73, 1849-1892.

[10] Marinacci, M. (2015): "Model uncertainty," *Journal of the European Economic Association*, 13, 998-1076.

[11] Savage, L.J. (1954): *The Foundations of Statistics.* New York: John Wiley and Sons.

[12] von Neumann J. and O. Morgenstern (1947): *The Theory of Games and Economic Behavior* (Second Ed.) Princeton: Princeton University Press.