http://selod.ensae.net/m1

Chapter 7 - Choice under uncertainty

Harris SELOD

Paris School of Economics (selod@ens.fr)

November 9, 2007

The issue

- Facing uncertainty, how do agents behave?
- What are their preferences over actions which results are not known in advance?
- What is the utility derived from an action with uncertain outcomes?

Random games and expected gains

Example 1: “Heads or tails?”

- Heads: $p = \frac{1}{2}$, $G = 100$
- Tails: $p = \frac{1}{2}$, $G = 0$
- How much (i.e what entrance right) would you agree to pay to play the game?

$$E(G) = \frac{1}{2}.100 + \frac{1}{2}.0 = 50$$

$$E(G) - r > 0 \Longrightarrow r < 50$$

Example 2: St-Petersburg game (paradox)

- Repeat "heads or tails?" until tails then stop.
- The player gets 2^n where n is the rank of the last throw. G denotes the gain.
- The expected value (expectation) of the gain is:

$$E(G) = \frac{1}{2}.2 + \frac{1}{4}.4 + ... + \frac{1}{2^n}.2^n + ... = \sum_{n=1}^{\infty} \frac{1}{2^n}.2^n = \infty$$

- Who would accept to pay an infinite entrance right to play this game? Not me!

Solving the paradox

- Substitute the utility of the gain for the gain. Then compute the expected utility of the gain instead of the expected gain:

$$u(x) = \log(x) \Longrightarrow E(u(G)) = \sum_{n=1}^{\infty} \frac{1}{2^n} \log 2^n \underset{n\ \infty}{\longrightarrow} 2\log 2$$

- A potential player would be ready to pay at most an entrance right r satisfying:

$$\underbrace{2\log 2}_{E(u(G))} - \underbrace{\log(r)}_{u(r)} = 0 \Longrightarrow r = 4 \quad (not + \infty!)$$

Lessons from theses examples

- The choice of u matters.
- Graph of expected utility (Example 1).

$$u \text{ concave} \Longrightarrow \underbrace{E(u(G))}_{\text{expected utility of the game}} < \underbrace{u(E(G))}_{\text{utility of getting E(G) for sure}}$$

- The concavity of u captures "risk aversion".
- When the player is "risk-averse", the maximum entrance right (his willingness to pay) must be smaller than the expected gain:

$$E(u(G)) = u(r) < u(E(G)) \Rightarrow r < E(G)$$

Questions:

- How do agents choose between random games?
  $\Rightarrow$ formalize risk

- In particular, what types of utility functions could be applied to actions with uncertain outcomes?
  $\Rightarrow$ **expected utility theory**

Formalizing the expected utility criterion

Notations

- Consider a set of N possible events.
- p_n is the objectively-known probability that event $n \in \{1, ..., N\}$ occurs.
- $\sum_{n=1}^{N} p_n = 1$ (complete system of events).

Definition (Lottery)

*A **simple lottery** L is a list* $L = (p_1, p_2, ..., p_N)$ *with* $p_n \geq 0$ *for all n and* $\sum_{n=1}^{N} p_n = 1$*, and where* p_n *is interpreted as the probability of outcome n occuring.*

Remark:

A lottery is defined over a set of uncertain outcomes.

It is a formal setting in which these uncertain outcomes are described by known probabilities.

The set of simple lotteries is a $(N-1)$-simplex denoted $\Delta(N-1) = \{p \in \mathbb{R}^N_+ : p_1 + p_2 + ...p_N = 1\}$

Definition (Compound lottery)

Consider K simple lotteries:
$L_k = (p_{1k}, p_{2k}, ..., p_{Nk})$ *for* $k = 1, ..., K$.

Consider $(\alpha_1, \alpha_2, ..., \alpha_K) \in \Delta(K-1)$ *where* α_k *is the probability of lottery k occurring* ($\alpha_k = Prob(L_k)$)

*The **compound lottery** L is the collection of the K lotteries and their probabilities of occurrence:*

$$L = (L_1, ..., L_K; \alpha_1, ..., \alpha_K)$$

Remark: The outcomes of the compound lottery are lotteries themselves.

Definition (Reduced lottery)

Any compound lottery determines a simple lottery $L = (p_1, ..., p_n, ..., p_N)$ *called a* ***reduced lotery***, *where*

$$p_n = \alpha_1 p_{n1} + ... + \alpha_k p_{nk} + ... + \alpha_K p_{nK}$$

with:

p_n *is the probability of event n*

α_k *is the probability of lottery k occurring*

p_{nk} *is the probability of event n in lottery k*

It is the simple lottery that generates the same ultimate distribution over outcomes as the compound lottery.

- Example 1
- Example 2
- Remark:

  No bijection between simple and compound lotteries

The preferences of agents over lotteries

- **Question**: what properties do we want for this type of preferences?
- We will make reasonable assumptions and build adequate utility functions. The issue is how to associate a utility to an uncertain event (lottery) which can have N outcomes.

- **Assumptions**

  $\succsim$ forms a complete and transitive (rational) binary relation on $\Delta(N-1)$

  Two compound lotteries that can be reduced to the same simple lottery will have to be considered equivalent ("consequentialist premise"; agents will only care about final results and not about the path towards these results).

- **The problem** is to represent $\succsim$ with a utility function with a specific structure: the **expected utility form.**

Definition (Expected utility form)

A utility function $U : \Delta(N-1) \rightarrow \mathbb{R}$ has an ***expected utility form*** *if there is an assignment of numbers $(u_1, ..., u_N)$ to the N outcomes, such that for any simple lottery*

$$L = (p_1, p_2, ..., p_N) \in \Delta(K-1)$$

we have:

$$U(L) = \sum_{n=1}^{N} p_n u_n = u_1 p_1 + ... + u_n p_n + ... + u_N p_N$$

A utility function $U : \Delta(N-1) \rightarrow \mathbb{R}$ with the expected utility form is called a ***von Neumann-Morgenstern (VNM) expected utility function****.*

Remarks

- It is a function applied to a lottery (a vector of probabilities).
- It is a linear combination of probabilities (the mathematical expectation of utilities, with u_k the utility associated with event k). Linearity may not be realistic because it confers a particular structure to U, implying that **the expected utility form is a cardinal property** of utility functions defined on the space of lotteries (it is not preserved under *any* strictly increasing transformation, but only under strictly increasing linear transformations).

Remarks (continued)

- There is sometimes an ambiguity in the literature on what is called the VNM function. According to Mas-Collel, Whinston and Green, **U(.) is the VNM function.** When $u_k = u(k)$, **u(.) is called the Bernoulli function.** Some authors call $u(.)$ the VNM function.

Proposition (Linearity of U)

U has the expected utility form

$\Longleftrightarrow$

U is linear: $\forall\, L_1, ..., L_K \in \Delta(N-1),\ \forall\, \alpha_1, ..., \alpha_K \in \Delta(N-1)$

$$U(\sum_{k=1}^{K} \alpha_k L_k) = \sum_{k=1}^{K} \alpha_k U(L_k)$$

Proof.

Remark: We are not "adding" lotteries but substituting them with probability α_k.

Axiomatic characterization (of preferences over risky alternatives)

Questions:

- What proporties do we want to impose to such preferences?
- When can we find a VNM utility function to represent them?

Two central axioms:

Definition (Continuity axiom)

$\forall$ L,L',L'' $\in \Delta(N-1)$, the sets
$\{\alpha \in [0,1] : \alpha L + (1-\alpha)L' \succsim L''\}$ *and*
$\{\alpha \in [0,1] : \alpha L + (1-\alpha)L' \precsim L''\}$ *are closed.*

Interpretation: small changes in probabilities do not change the ordering between two lotteries.

Remark: we are not "adding" L' to L but substituting it to L (with a probability $1-\alpha$)

Example (re-stated from Mas-Collel p.171):

- Consider three events:
  trip by car / death by accident / staying home

  $L = (1, 0, 0)$ is a safe trip by car with probability 1

  $L'' = (0, 0, 1)$ is staying home with probability 1

  Assume $L \succsim L''$

  $L' = (0, 1, 0)$ is death by accident with probability 1

Example (continued):

- If the continuity axiom is satisfied then a trip by car even with a small chance of death by accident ($\alpha L + (1-\alpha)L'$ *with* α *close to* 1) is preferred over staying home (L'').
- There is no discontinuity (of the type "safety first") as with lexicographic preferences.

Definition (Independence axiom)

$\forall\ L, L', L'' \in \Delta(N-1), \forall\ \alpha \in [0,1]$, *we have:*

$$L \succsim L' \iff \alpha L + (1-\alpha)L'' \succsim \alpha L' + (1-\alpha)L''$$

Interpretation: If $L \succsim L'$, a compound lottery must not change the preference order: considering the possibility of substituting a third lottery with the same probability $(1-\alpha)$ does not change the order between any two lotteries. Which lottery is preferred is independent of the third lottery.

Remark: This is a fundamental (but controversial) axiom in the theory of choice under uncertainty.

Discussion of the independence axiom:

- Experiments show that the independence axiom is sometimes violated.
- Examples:

  the Machina paradox

  the Allais paradox

Representation of indifference curves

- 3-dimension representation. **Graph.**
- 2-dimension representation. **Graph.**

Proposition

A complete and transitive $\succsim$ satisfying the continuity and independence axioms has ***linear and parallel indifference curves****.*

Proof.

The expected utility theorem

Proposition (Expected utility theorem)

Consider $\succsim$ a complete and transitive preference relation on $\Delta(N-1)$ satisfying the continuity and independence axioms

$\Longrightarrow$

$\succsim$ can be represented by a utility function having the expected utility form.

$$\text{(i.e. } \underbrace{L}_{(p_1,p_2,\ldots,p_N)} \succsim \underbrace{L'}_{(p'_1,p'_2,\ldots,p'_N)} \Longleftrightarrow \underbrace{U(L)}_{\sum_{n=1}^{N} u_n p_n} \geq \underbrace{U(L')}_{\sum_{n=1}^{N} u_n p'_n}\text{)}$$

Measuring risk aversion (monetary lotteries)

Notations

- The context is now continuous (i.e. the general case; useful for monetary lotteries).
- Consider a continuous real-valued random variable X (for instance income). The lottery is now defined by X's probability density function (p.d.f. denoted f) or cumulative distribution function (c.d.f. denoted F). Recall that:

$$F(x) = P(X \leq x) = \int_{-\infty}^{x} f(t)dt = \int_{-\infty}^{x} dF(t)$$

- $\mathfrak{L}$ the set of lotteries is the set of c.d.f. on $[a, +\infty[$

Definitions

Definition (Expected utility form)

U has an expected utility form if:

$\exists\, u : \mathbb{R} \rightarrow \mathbb{R}$ *($u(x)$ being the utility derived from income x) such that $\succsim$ is represented by:*

$U : \mathfrak{L} \rightarrow \mathbb{R}$
$U(F) = \int u(x) dF(x)$

Definition (Risk aversion)

*An individual is **risk averse** if he prefers, $\forall\, F \in \mathfrak{L}$, the degenerate lottery denoted F_M that yields $\int x dF(x)$ with certainty.*

Interpretation: the individual prefers gaining the expected value with certainty than incurring some risk but gaining the same value on average.

Definition (Risk neutrality)

*An agent is **risk neutral** if* $F \sim F_M$.

Proposition

If U is VNM,

$$\textit{risk aversion} \iff \underbrace{\int u(x)dF(x)}_{E(u(X))\ \textit{risky}} \leq \underbrace{u(\int xdF(x))}_{u(E(X))} \ \forall\ F \in \mathfrak{L}$$

This is Jensen's inequality which defines concave functions. Hence **risk aversion** ⇔ **concavity of u(.)**

Decreasing marginal utility of income. The utility gain from an extra euro is lower than the utility loss of having a euro less. Hence risk aversion = the fear of losing.

Consider $u(.)$ a Bernoulli function and $F \in \mathfrak{L}$

Definition (Certainty equivalent)

*The **certainty equivalent** of F is the amount $ce(F, u)$ such that*

$$u(ce(F, u)) = \int u(x)dF(x)$$

It is the amount of money that, if gained with certainty, provides the same utility as the gamble F.

This is the willingness to pay for F.
Cf. the entrance right in our examples.

Definition (Risk premium)

*The **risk premium** associated with F is the difference:*

$$\rho = \int x dF(x) - ce(F, u)$$

$$ce(F, u) = \int x dF(x) - \rho$$

It is the loss of income that can be conceded in order to get rid of the risk (and obtain the certainty equivalent).

It measures the gap between the expected value of the gamble and the certainty equivalent. It is “positively correlated” with risk aversion.

Measures of risk aversion Consider $u(.)$ a Bernoulli function.

We can define:

Definition (Absolute risk aversion coefficient)

The absolute risk aversion coefficient (also called the ***Arrow-Pratt coefficient of absolute risk aversion****) is:*

$$r_A(x, u) = -\frac{u''(x)}{u'(x)} > 0$$

It is a concavity index for $u(.)$ which is invariant to positive linear transformations of $u(.)$.

$u(.)$ can be recovered from r_A.

Example of a **CARA (Constant Absolute Risk Aversion) utility function**:

$$u(x) = -e^{-cx} \text{ with } c > 0$$

Definition (Relative risk aversion coefficient)

*The **relative risk aversion coefficient**:*

$$r_R(x,u) = -x\frac{u''(x)}{u'(x)} = x.r_A(x,u)$$

Example of a **CRRA (Constant Relative Risk Aversion) utility function**:

$$u(x) = \frac{x^{1-\gamma}}{1-\gamma} \text{ with } \gamma > 0$$