Lecture notes 5: Uncertainty

Choice under uncertainty

- Uncertainty an unsepararable part of individual decision making
- But what is uncertainty?
- This far, there has no need to distinguish **actions** and **consequences**
- Undcertainty does, by definition, separate the two: the correspondence between actions and consequences is going to be *stochastic*
- How should a "rational" decision maker value different **lotteries** over consequences?
- Not an obvious thing to do

Example (St Petersburg paradox, Bernoulli 1738)

How much would you pay for a gamble where one tosses a fair coin until one wins and after each toss, the prize money doubles?

- The paradox: one is usually(!) willing to pay only limited price for the lottery whereas the expected monetary value is

$$2 \cdot \frac{1}{2} + 2^2 \cdot \frac{1}{2^2} + 2^3 \cdot \frac{1}{2^3} + ... = 1 + 1 + 1 + ... = \infty$$

- ...hence one has to maximize the expected utility value of the money
- For example, with utility function $u(x) = \log_2 x$, then the value of the gamble is 1

Uncertainty

- The aim here is develop systematically a theory that allows us to assess the value of different lotteries
- Let there be a set X (finite, for simplicity) of consequences or outcomes
- Uncertainty is reflected by a **lottery** p over X, which is just a probability distribution such that

$$p(x) \geq 0, \text{ for all } x \in X \quad \text{and} \quad \sum_{x \in X} p(x) = 1$$

- For any two lotteries p and q, and parameter $\lambda \in [0,1]$, denote by $\lambda \cdot p + (1-\lambda) \cdot q$ the mixed lottery such that

$$(\lambda p + (1-\lambda)q)(x) = \lambda \cdot p(x) + (1-\lambda)q(x), \text{ for all } x \in X$$

- Denote the degenerate lottery that puts all the probablity mass on the consequence x by 1_x
- The decisions then concern feasible lotteries, and hence the preference relation should be defined over the **lottery space**

$$L = \left\{ p \in \mathbb{R}_+^{|X|} : \sum_{x \in X} p(x) = 1 \right\}$$

- The observable choices $\succsim$ are now defined over L, i.e. $\succsim \subset L \times L$

Expected utility maximization

- Rationality implies that the decision maker's preferences $\succsim$ are **complete** and **transitive**
- In particular, the agent has well defined preferences over the **pure** outcomes $\{1_x : x \in X\}$
- But all complete and transitive preferences plausible?

Examples

(bad examples)

- Preference for certainty: $p \succsim q$ iff $\max_x p(x) \geq \max_x q(x)$
  - But this is independent of the consequences!
- Avoidance of the worst case: $p \succsim q$ iff $\min_x \{1_x : p(x) > 0\} \succsim \min_x \{1_x : q(x) > 0\}$
  - But this is very sensitive to changes in probabilities and ignores what happens to the probablilities of outcomes other than the worst case!
- Comparing most likely consequences: $p \succsim q$ iff $1_{x(p)} \succsim 1_{x(q)}$ where $x(p) = \max_x p(x)$ and $x(q) = \max_x q(x)$
  - But this is again very sensitive to changes in probabilities and ignores what happens to the probabilities of the outcomes other than $x(p)$ and $x(q)$!

- Objective: to impose conditions on preferences that avoid the obvious pitfalls in evaluating lotteries

NM1 (Continuity) For all $p, q, r \in L$, if $p \succ q \succ r$, then there are $\lambda, \mu \in (0, 1)$ such that
$\lambda \cdot p + (1 - \lambda) \cdot r \succ q \succ \mu \cdot p + (1 - \mu) \cdot r$

- Equivalently, the upper and lower contour sets of $\succsim$ are closed
- Continuity axiom is occasionally called the "Archimedean axiom"

- The following condition implies that a choice between two lotteries should be independent from the probability that this comparison is actually conducted
- If two lotteries are changed in the same way, their relative ranking should not change

NM2 (Independence) For all $p, q, r \in L$ and $\lambda \in (0,1)$, if $p \succsim q$ then $\lambda \cdot p + (1-\lambda) \cdot r \succsim \lambda \cdot q + (1-\lambda) \cdot r$

- A direct implication of independence is that if $p \sim q$, then $\lambda \cdot p + (1-\lambda) \cdot r \sim \lambda \cdot q + (1-\lambda) \cdot r$, for any $\lambda \in (0,1)$

Lemma

If a preference ordering $\succsim$ satisfies independence (NM2), then

$$\begin{aligned} p &\succsim q \quad \textit{if and only if} \\ \lambda \cdot p + (1-\lambda) \cdot q &\succsim \lambda' \cdot p + (1-\lambda') \cdot q, \quad \textit{for all } \lambda \geq \lambda' \end{aligned}$$

Proof.

By independence, for any $\lambda \in (0,1)$

$$\lambda \cdot p + (1-\lambda) \cdot q \succsim \lambda \cdot q + (1-\lambda) \cdot q = q.$$

Applying this with respect to $\mu \in (0,1)$,

$$\begin{aligned} \lambda \cdot p + (1-\lambda) \cdot q &\succsim \mu \cdot (\lambda \cdot p + (1-\lambda) \cdot q) + (1-\mu) \cdot q \\ &= \mu\lambda \cdot p + (1-\mu\lambda) \cdot q. \end{aligned}$$

By choosing $\mu = \lambda'/\lambda$, the result ensues. $\square$

Lemma

If a preference ordering $\succsim$ satisfies continuity (NM1), and independence (NM2), then for any $p \succ q \succ r$ there is a unique $\alpha \in (0,1)$ such that $\alpha \cdot p + (1-\alpha) \cdot r \sim q$

Proof.

Consider sets $\{\lambda \in [0,1] : \alpha \cdot p + (1-\alpha) \cdot r \succ q\}$ and $\{\lambda \in [0,1] : q \succ \alpha \cdot p + (1-\alpha) \cdot r\}$. By construction, the sets are disjoint. By the previous lemma, they are intervals and, by continuity, open. Hence they do not cover $[0,1]$, and there is an element α of $[0,1]$ not in $\{\lambda \in [0,1] : \alpha \cdot p + (1-\alpha) \cdot r \succ q\}$ or $\{\lambda \in [0,1] : q \succ \alpha \cdot p + (1-\alpha) \cdot r\}$. By construction, $\alpha \cdot p + (1-\alpha) \cdot r \sim q$. By the previous lemma, α is unique. $\square$

Theorem

(von Neumann and Morgenstern 1944) Let X be a finite set. Then If a preference ordering $\succsim$ satisfies NM1-NM2 if and only if there is a function $u : X \to \mathbb{R}$ such that

$$\sum_{x \in X} p(x)u(x) \geq \sum_{x \in X} q(x)u(x) \quad \text{if and only if} \quad p \succsim q.$$

Moreover, the function u is unique up to positive linear transformation

Proof.

Identify x^* and x_* such that

$$1_{x^*} \succsim 1_x \succsim 1_{x*}, \quad \text{for all } x \in X$$

By Lemmata 4 and 3 there is, for any $x \in X$, a unique α_x such that

$$1_x \sim \alpha_x \cdot 1_{x^*} + (1 - \alpha_x) \cdot 1_{x_*}$$

Any lottery p can be written as a mixture

$$p = p(x) \cdot 1_x + (1 - p(x)) \cdot p^x$$

where

$$p^x(y) = \frac{p(y)}{1 - p(x)}, \quad \text{for all } y \in X \backslash \{x\}$$

$\square$

Proof.

[cont.] By independence,

$$p \sim p(x)[\alpha_x \cdot 1_{x^*} + (1-\alpha_x) \cdot 1_{x_*}] + (1-p(x)) \cdot p^x$$

By induction on the cardinality of X,

$$p \sim \sum_x p(x)[\alpha_x \cdot 1_{x^*} + (1-\alpha_x) \cdot 1_{x_*}]$$

or, equivalently,

$$p \sim \left(\sum_x p(x)\alpha_x\right) \cdot 1_{x^*} + \left(1 - \sum_x p(x)\alpha_x\right) \cdot 1_{x_*}$$

By choosing $u(x) = \alpha_x$ for all $x \in X$, and by Lemma 3,

$$p \succsim q \quad \text{if and only if} \quad \sum_x p(x)u(x) \geq \sum_x q(x)u(x).$$

$\square$

Proof.

[cont.] To see the uniqueness, note that for any representation v of preferences $\succsim$,

$$v(x) = \alpha_x v(x^*) + (1 - \alpha_x) v(x_*), \text{ for all } x \in X$$

Find $a \in \mathbb{R}_{++}$ and $b \in \mathbb{R}$ such that

$$v(x^*) = au(x^*) + b \quad \text{and} \quad v(x_*) = au(x_*) + b$$

Since also

$$u(x) = \alpha_x u(x^*) + (1 - \alpha_x) u(x_*), \text{ for all } x \in X$$

it follows that

$$\begin{aligned} v(x) &= \alpha_x[au(x^*) + b] + (1 - \alpha_x)[au(x_*) + b] \\ &= a[\alpha_x u(x^*) + (1 - \alpha_x) u(x_*)] + b \\ &= au(x) + b \end{aligned}$$

- Provides a definition of a utility function (sometimes called Bernoulli utility function); it represents preferences over lotteries
- Now also the intensities, or the curvature of u, matters
- However, again, no conclusion about "real" utilities should be drawn, there are infinitely many representations
- Gives a justification for the expected utility maximization (rather than, say, median)
- Normative argument: if the axioms are accepted, there has to be a (Bernoulli) utility function

- In the proof we used an induction argument
- Without finiteness of X, the indictive step requires further assumption
- The expected utility theory is very useful for modeling purposes and for normative argumentation
- However, it does less well in experiments

Example (Kahnemann and Tversky 1979)

(orignally due to Allais, 1953) There are two choice scenarios:

1. Choice between lotteries
   1. $0.33 \cdot 2500 + 0.66 \cdot 2400 + 0.01 \cdot 0$
   2. $1 \cdot 2400$
2. Choice between lotteries
   1. $0.33 \cdot 2500 + 0.67 \cdot 0$
   2. $0.34 \cdot 2400 + 0.66 \cdot 0$

- Of the subjects, 82% chose 1a and 83% chose 2a which means that at least 65% chose both 1b and 2a

Example (cont.)

However, there is no utility function u that is consistent with the choices and expected utility maximization:

$$0.33 \cdot u(2500) + 0.66 \cdot u(2400) + 0.01 \cdot u(0) < u(2400)$$

and

$$0.33 \cdot u(2500) + 0.67 \cdot u(0) > 0.34 \cdot u(2400) + 0.66 \cdot u(0)$$

i.e.

$$0.66 \cdot (u(2400) - u(0)) < 0.66 \cdot (u(2400) - u(0))$$

Application: Risk Aversion

- Assume that $X = \mathbb{R}_+$, and interpret x as "money"
- Let vNM preferences $\succsim$ over simple lotteries L (with finite support, for simplicity) be represented by a utility function $u : \mathbb{R}_+ \to \mathbb{R}$
- What are plausible assumptions concerning u?
- More money is strictly better $u(x) > u(y)$ if $x > y$

Risk

- From now on, consequences are monetary amounts: $x \in \mathbb{R}_+$ is the final wealth of the decision maker
- Analyze different $u : \mathbb{R}_+ \to \mathbb{R}$
- p denotes the distribution function of a monetary lottery

- By the expected utility theorem:

$$U(p) = \sum_{x \geq 0} u(x)p(x),$$

- When do the integrals and sums above converge? (cf. St.Petersburg's paradox)

- When considering expected utility, we can consider variations in each of these two components.
  1. Risk attitudes: Fix p and compare different u
  2. Riskiness of lotteries: Fix u and compare different p

Definition

The **certainty equivalent** $c(p,u)$ of a lottery p for a decision maker with utility function u is defined by

$$u(c(p,u)) = \sum_x u(x)p(x)$$

- We can discuss attitudes towards risk by comparing the certainty equivalents of a fixed lottery under different utility functions

Definition

A decision maker with a utility function u is **risk averse** if, for all p,

$$c(p, u) \leq \sum_x xp(x)$$

Proposition

Utility function u is **risk averse** *if and only if it is concave (exercise, use Jensen's inequality)*

- Risk loving attitudes are defined with the opposite inequalities

- Can second derivatives be used to measure risk aversion? But not robust against linear transformations of u!

Definition

The Arrow-Pratt measure of absolute risk aversion, $r_A(x, u)$ of utility function u at wealth level x is given by:

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

- $r_A(x, u)$ reflects the curvature of u at x without accounting linear transformations
- The following results shows that $r_A(x, u)$ is a good measure of risk aversion

- When is u_1 unambigiously more risk averse than u_2?

Proposition

The following are equivalent:
i) $r_A(x, u_2) \geq r_A(x, u_1)$ *for all* x
ii) $c(p, u_2) \leq c(p, u_1)$ *for all* p
iii) There is a concave function $\phi(\cdot)$ *such that* $u_2(x) = \phi(u_1(x))$

- The more-risk-averse-than relation is a partial ordering of utility functions: it is not possible compare all u and u' in this sense

- Do wealthier individuals take bigger risks than others?

Definition

u exhibits **decreasing** (constant) **absolute risk aversion**, DARA, (CARA, resp.) if $r_A(x, u)$ is a decreasing (constant, resp.) function of x

- If u exhibits DARA, then the decion maker is less sensitive towards risk when his wealth increases, i.e. he is willing to pay less to get rid of risk as his wealth increases
- The family of utility functions exhibiting CARA is given by the following conditions: there is a constant λ such that $\lambda = -u''(x)/u'(x)$ for all x, and there are $a > 0$ and b such that for all

$$u(x) = -ae^{-\lambda x} + b$$

- If u exhibits DARA, then , for any $x_2 < x_1$ there is a concave function $\phi(\cdot)$ such that $\phi(u(x_1 + z)) = u(x_2 + z)$ for all z
- Alternatively, if $u_x(z) = u(x + z)$ for all z, then $x - c(p, u_x)$ is decreasing in x, for any p
- Thus a person with DARA utility function is less risk-averse when richer

- A stronger condition: u exhibits **decreasing** (constant) **relative risk aversion**, DRRA, (CRRA, resp.) if $r_R(x,u) := xr_A(x,u)$ decreases in x
- Requires decreasing risk aversion subject to proportional gambles: e.g. willingness to risk all ones wealth to double it
- Strong condition, DRRA implies DARA: $r_A'(x,u) < -r_A(x,u)/x$ for all x
- Plausible and often verified empirically
- But: Rabin's paradox

Example (Rabin's paradox)

Assume weak risk aversion, e.g. DARA with

$$\frac{1}{2}u(x-10)+\frac{1}{2}u(x+11)\leq u(x),\ \text{for all } x\in\mathbb{R}_+$$

i.e. gamble of losing 10€ and winning 11€ with equal probability is weakly rejected at all wealth levels. Often observed in calibration excercises. Can this be a general phenomenon?

Example

If yes, then, for all wealth levels x

$$u(x-10)-u(x) \leq u(x+11)-u(x)$$

or

$$\frac{u(x-10)-u(x)}{10} \leq \frac{10}{11} \cdot \frac{u(x+11)-u(x)}{11}.$$

Thus the marginal utility u' drops by at least $10/11$ in every additional 21€.

Example

- But then, adding 1000€ to the DM's wealth would mean that the decision maker"s marginal value of a € drops by proportion

$$\left(\frac{10}{11}\right)^{1000/21} = 0.012.$$

Moreover, there is **no** compensating prize M that would render acceptable a gamble, where one loses 100€ with probability $1/2$ and wins prize M with probability $1/2$!

$$\sum_{t=1}^{5}\left(\frac{10}{11}\right)^{t} \leq \sum_{t=6}^{T}\left(\frac{10}{11}\right)^{t}, \quad \text{for all } T$$

- Hence global risk aversion is not a plausible story
- The importance of reference dependence

- How do changes in the distribution affect expected payoff?
- Distribution p **first order stochastically dominates** q if $\sum_{x \leq y} p(x) \leq \sum_{x \leq y} q(x)$ for all y
- This is equivalent to saying that, for all nondecreasing functions u,
$$\sum_x u(x)q(x) \geq \sum_x u(x)p(x)$$
- Thus if p first order stochastically dominates q, the expected value from p is higher, i.e. shifting probability mass forward increases expected payoff

- But first order stochastic dominance may concern the mean values of the lotteries and hence be independent of their risk properties
- Distribution p **second order stochastically dominates** q if they have the same mean

$$\sum_x xp(x) = \sum_x xq(x),$$

and for all $y \geq 0$ we have

$$\sum_{x \geq y} p(x) \geq \sum_{x \geq y} q(x)$$

- Shift from p to q is a **mean preserving spread**, disliked by any risk-averter