14.12 Game Theory Lecture Notes
Theory of Choice

(Lecture 2)

1 The basic theory of choice

We consider a set X of alternatives. Alternatives are mutually exclusive in the sense that one cannot choose two distinct alternatives at the same time. We also take the set of feasible alternatives exhaustive so that a player's choices will always be defined. Note that this is a matter of modeling. For instance, if we have options Coffee and Tea, we define alternatives as $C =$ Coffee but no Tea, $T =$ Tea but no Coffee, $CT =$ Coffee and Tea, and $NT =$ no Coffee and no Tea.

Take a relation $\succeq$ on X. Note that a relation on X is a subset of $X \times X$. A relation $\succeq$ is said to be *complete* if and only if, given any $x, y \in X$, either $x \succeq y$ or $y \succeq x$. A relation $\succeq$ is said to be *transitive* if and only if, given any $x, y, z \in X$,

$$[x \succeq y \text{ and } y \succeq z] \Rightarrow x \succeq z.$$

A relation is a *preference relation* if and only if it is complete and transitive. Given any preference relation $\succeq$, we can define strict preference $\succ$ by

$$x \succ y \iff [x \succeq y \text{ and } y \not\succeq x],$$

and the indifference $\sim$ by

$$x \sim y \iff [x \succeq y \text{ and } y \succeq x].$$

A preference relation can be *represented* by a utility function $u : X \to \mathbb{R}$ in the following sense:

$$x \succeq y \iff u(x) \geq u(y) \qquad \forall x, y \in X.$$

The following theorem states further that a relation needs to be a preference relation in order to be represented by a utility function.

Theorem 1 *Let X be finite. A relation can be presented by a utility function if and only if it is complete and transitive. Moreover, if $u : X \to \mathbb{R}$ represents $\succeq$, and if $f : \mathbb{R} \to \mathbb{R}$ is a strictly increasing function, then $f \circ u$ also represents $\succeq$.*

By the last statement, we call such utility functions ordinal.

In order to use this ordinal theory of choice, we should know the agent's preferences on the alternatives. As we have seen in the previous lecture, in game theory, a player chooses between his strategies; and his preferences on his strategies depend on the strategies played by the other players. In order to apply this theory to games directly, we might need to restrict ourselves to the cases where each player knows which strategies the other players play. This is clearly too restrictive, hence we need a theory of decision-making under uncertainty.

2 Decision-making under uncertainty

We consider a finite set Z of prizes, and the set P of all probability distributions $p : Z \to [0, 1]$ on Z, where $\sum_{z \in Z} p(z) = 1$. We call these probability distributions lotteries. A lottery can be depicted by a tree. For example, in Figure 1, Lottery 1 depicts a situation in which if head the player gets \$10, and if tail, he gets \$0.

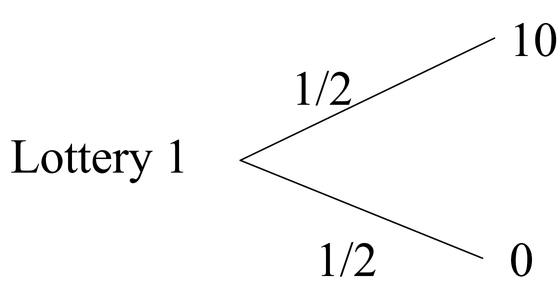

Figure 1:

In game theory and more broadly when agents make their decision under uncertainty, we do not have the lotteries as in casinos where the probabilities are generated by some machines and as we have defined above where the probabilities are given. It has been shown by Savage (1954) under certain conditions that a player's beliefs can be represented by a (unique) probability distribution. Using these probabilities, we can represent our acts by lotteries.

We would like to have a theory that constructs a player's preferences on the lotteries from his preferences on the prizes. The most well-known such theory is the theory of expected utility maximization by Von Neumann and Morgenstern. A preference relation $\succeq$ on P is said to be represented by a von Neumann-Morgenstern utility function $u : Z \to \mathbb{R}$ if and only if

$$p \succeq q \iff U(p) \equiv \sum_{z \in Z} u(z)p(z) \geq \sum_{z \in Z} u(z)q(z) \equiv U(q) \tag{1}$$

for each $p, q \in P$. Note that $U : P \to \mathbb{R}$ represents $\succeq$ in ordinal sense. That is, the agent acts as if he wants to maximize the expected value of u. For instance, the expected utility of Lottery 1 for our agent is $E(u(\text{Lottery } 1)) = \frac{1}{2}u(10) + \frac{1}{2}u(0)$.[1]

The necessary and sufficient conditions for a representation as in (1) are as follows:

Axiom 1 $\succeq$ *is complete and transitive.*

This is necessary by Theorem 1, for U represents $\succeq$ in ordinal sense. The second condition is called *independence* axiom, stating that a player's preference between two lotteries p and q does not change if we toss a coin and give him a fixed lottery r if "tail" comes up.

Axiom 2 *For any $p, q, r \in P$, and any $a \in (0, 1]$, $ap + (1-a)r \succ aq + (1-a)r \iff p \succ q$.*

Let p and q be the lotteries depicted in Figure A. Then, the lotteries $ap + (1-a)r$ and $aq + (1-a)r$ can be depicted as in Figure B, where we toss a coin between a fixed lottery r and our lotteries p and q. Axiom 2 stipulates that the agent would not change his mind after the coin toss. Therefore, our axiom can be taken as an axiom of "dynamic consistancy" in this sense.

The third condition is purely technical, and called *continuity* axiom. It states that there are no "infinitely good" or "infinitely bad" prizes.

[1] If Z were a continuum, like $\mathbb{R}$, we would compute the expected utility of p by $\int u(z)p(z)dz$.

Axiom 3 *For any $p, q, r \in P$, if $p \succ r$, then there exist $a, b \in (0,1)$ such that $ap + (1 - a)r \succ q \succ bp + (1 - r)r$.*

Axioms 2 and 3 imply that, given any $p, q, r \in P$ and any $a \in [0, 1]$,

$$\text{if } p \sim q, \text{ then } ap + (1 - a)\, r \sim aq + (1 - a)r. \tag{2}$$

This has two implications:

1. The indifference curves on the lotteries are straight lines.

2. The indifference curves, which are straight lines, are parallel to each other.

To illustrate these facts, consider three prizes z_0, z_1, and z_2, where $z_2 \succ z_1 \succ z_0$. A lottery p can be depicted on a plane by taking $p(z_1)$ as the first coordinate (on the horizontal axis), and $p(z_2)$ as the second coordinate (on the vertical axis). $p(z_0)$ is $1 - p(z_1) - p(z_2)$. [See Figure C for the illustration.] Given any two lotteries p and q, the convex combinations $ap + (1 - a)\, q$ with $a \in [0, 1]$ form the line segment connecting p to q. Now, taking $r = q$, we can deduce from (2) that, if $p \sim q$, then $ap + (1 - a)\, q \sim aq + (1 - a)q = q$ for each $a \in [0, 1]$. That this, the line segment connecting p to q is an indifference curve. Moreover, if the lines l and l' are parallel, then $\alpha/\beta = |q'| / |q|$, where $|q|$ and $|q'|$ are the distances of q and q' to the origin, respectively. Hence, taking $a = \alpha/\beta$, we compute that $p' = ap + (1 - a)\, \delta_{z_0}$ and $q' = aq + (1 - a)\, \delta_{z_0}$, where δ_{z_0} is the lottery at the origin, and gives z_0 with probability 1. Therefore, by (2), if l is an indifference curve, l' is also an indifference curve, showing that the indifference curves are parallel.

Line l can be defined by equation $u_1 p(z_1) + u_2 p(z_2) = c$ for some $u_1, u_2, c \in \mathbb{R}$. Since l' is parallel to l, l' can also be defined by equation $u_1 p(z_1) + u_2 p(z_2) = c'$ for some c'. Since the indifference curves are defined by equality $u_1 p(z_1) + u_2 p(z_2) = c$ for various values of c, the preferences are represented by

$$\begin{aligned} U(p) &= 0 + u_1 p(z_1) + u_2 p(z_2) \\ &\equiv u(z_0)p(z_0) + u(z_1)p(z_1) + u(z_2)p(z_2), \end{aligned}$$

where

$$
\begin{aligned}
u(z_0) &= 0, \\
u(z_1) &= u_1, \\
u(z_2) &= u_2,
\end{aligned}
$$

giving the desired representation.

This is true in general, as stated in the next theorem:

Theorem 2 *A relation $\succeq$ on P can be represented by a von Neumann-Morgenstern utility function $u : Z \to R$ as in (1) if and only if $\succeq$ satisfies Axioms 1-3. Moreover, u and $\tilde{u}$ represent the same preference relation if and only if $\tilde{u} = au + b$ for some $a > 0$ and $b \in \mathbb{R}$.*

By the last statement in our theorem, this representation is "unique up to affine transformations". That is, an agent's preferences do not change when we change his von Neumann-Morgenstern (VNM) utility function by multiplying it with a positive number, or adding a constant to it; but they do change when we transform it through a non-linear transformation. In this sense, this representation is "cardinal". Recall that, in ordinal representation, the preferences wouldn't change even if the transformation were non-linear, so long as it was increasing. For instance, under certainty, $v = \sqrt{u}$ and u would represent the same preference relation, while (when there is uncertainty) the VNM utility function $v = \sqrt{u}$ represents a very different set of preferences on the lotteries than those are represented by u. Because, in cardinal representation, the curvature of the function also matters, measuring the agent's attitudes towards risk.

3 Attitudes Towards Risk

Suppose individual A has utility function u_A. How do we determine whether he dislikes risk or not, and whether he has a higher tolerence to risk than another individual B with utility function u_B?

The answer lies in the cardinality of the function u.

Let us first define a *fair* gamble, as a lottery that has expected value equal to 0. For instance, lottery 2

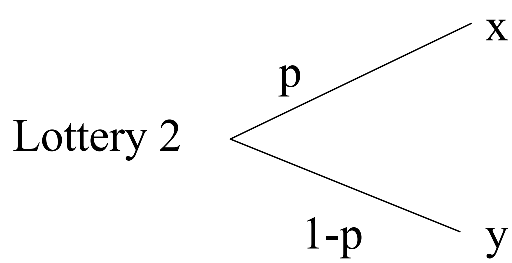

is a fair gamble if and only if $px + (1-p)y = 0$.

We define an agent as *Risk-Neutral* if and only if he is indifferent between accepting and rejecting all fair gambles. Thus, an agent with utility function u is risk neutral if and only if

$$E(u(\text{lottery 2})) = pu(x) + (1-p)u(y) = u(0)$$

for all p, x, and y.

This can only be true for all p, x, and y, if and only if the agent is maximizing the expected value, that is, $u(x) = ax + b$. Therefore, we need the utility function to be linear.

Therefore, an agent is risk-neutral if and only if he has a linear Von-Neumann-Morgenstern utility function.

An agent is *risk-averse* if and only if he rejects *all* fair gambles:

$$E(u(\text{lottery 2})) < u(0)$$
$$pu(x) + (1-p)u(y) < u(px + (1-p)y) \equiv u(0)$$

Now, recall that a function $g(\cdot)$ is strictly concave if and only if we have

$$g(\lambda x + (1-\lambda)y) > \lambda g(x) + (1-\lambda)g(y)$$

for all $\lambda \in (0,1)$. Therefore, risk-aversion is equivalent to having a strictly concave utility function.

Similarly, an agent is said to be risk seeking iff he has a strictly convex utility function.

Consider Figure 2. The cord AB is the utility difference that this risk-averse agent would lose by taking the gamble that gives W_1 with probablity p and W_2 with probability

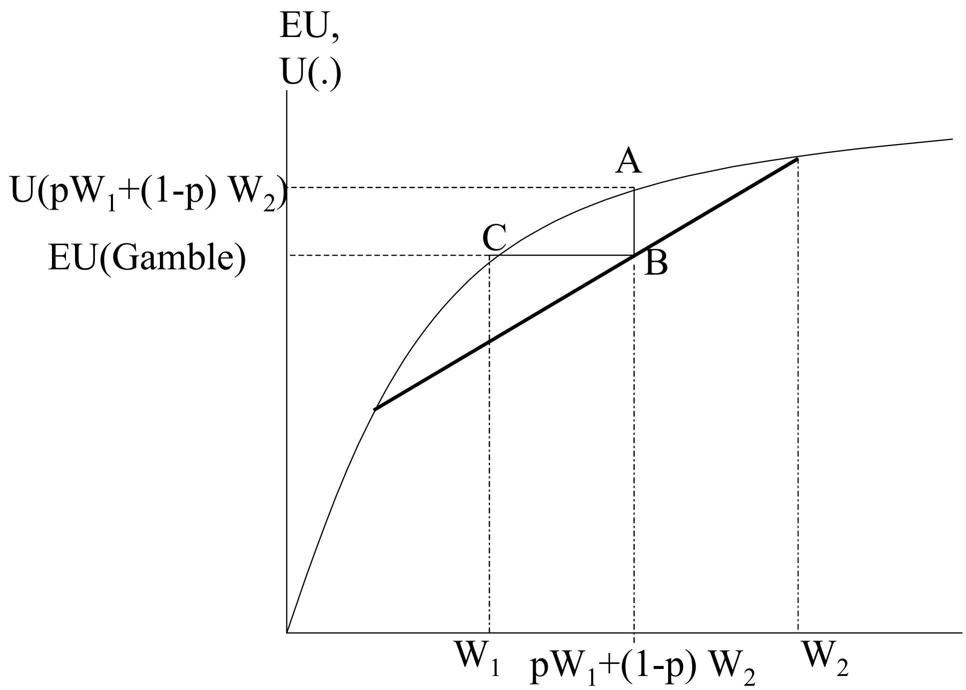

Figure 2:

$1-p$. BC is the maximum amount that she would pay in order to avoid to take the gamble. Suppose W_2 is her wealth level and $W_2 - W_1$ is the value of her house and p is the probability that the house burns down. Thus in absense of fire insurance this individual will have utility given by EU(gamble), which is lower than the utility of the expected value of the gamble.

3.1 Risk sharing

Consider an agent with utility function $u : x \mapsto \sqrt{x}$. He has a (risky) asset that gives \$100 with probability 1/2 and gives \$0 with probability 1/2. The expected utility of our agent from this asset is $EU_0 = \frac{1}{2}\sqrt{0} + \frac{1}{2}\sqrt{100} = 5$. Now consider another agent who is identical to our agent, in the sense that he has the same utility function and an identical asset, where two assets are statistically independent from each other. Imagine that our agents from a mutual fund by pooling their assets, each agent owning half of the mutual fund. This mutual fund gives \$200 with probability 1/4 (when both assets yield high dividends), \$100 with probability 1/2 (when only one on the assets gives high dividend), and gives \$0 with probability 1/4 (when both assets yield low dividends). Thus, each agent's share in the mutual fund yields \$100 with probability 1/4, \$50 with probability 1/2, and \$0 with probability 1/4. Therefore, his expected utility from the share in this mutual fund is $EU_S = \frac{1}{4}\sqrt{100} + \frac{1}{2}\sqrt{50} + \frac{1}{4}\sqrt{0} = 6.0355$. This is clearly larger than his expected utility from his own asset, therefore our agents gain by sharing the risk in their assets.

4 Measuring Risk-Aversion

If U_A is more 'curved' or more 'concave', then this will correspond to a more risk averse utility function or individual.

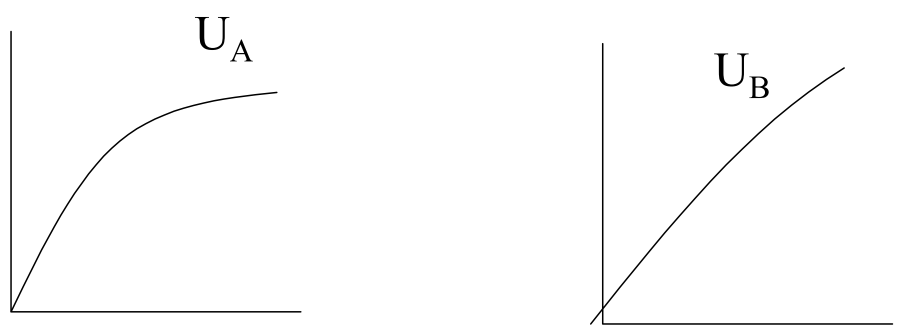

Figure 3:

How to formulize this?

U' is related to how concave a function is but since positive linear transforamtions are allowed its numerical value is of no interest. What matters is U'' relative to U'.

This is the Arrow-Pratt measures of risk-aversion:

Relative Degree of Risk-Aversion: $-\frac{WU''(W)}{U'(W)}$

Absolute Degree of Risk-Aversion: $-\frac{U''(W)}{U'(W)}$

A higher degree of risk aversion corresponds to a more risk-averse agent. More generally, A is more risk averse than B if and only if

- U_A is a concave transformation of U_B

- A's degree of absolute and relative risk-aversion is always higher than B's.

The rationale behind the relative risk aversion $(-WU''(W)/U'(W))$ is that an agent may get more (or less) risk averse as he gets richer, which is called the wealth effect. An agent's risk aversion does not change as his wealth change when his absolute risk-aversion $(-U''(W)/U'(W))$ is constant. In that case, the agent's utility function is defined by

$$u(x) = -e^{-\alpha x}$$

for some $\alpha > 0$, which is absolute risk aversion.