

Risk and Uncertainty 1

Mark Dean

GR6211 - Microeconomic Analysis 1

- Up until now, we have thought of people choosing between objects
 - Used cars
 - Hamburgers
 - Monetary amounts
- However, often the outcome of the choices that we make are not known
 - You are deciding whether or not to buy a share in AIG
 - You are deciding whether or not to put your student loan on black at the roulette table
 - You are deciding whether or not to buy a house that straddles the San Andreas fault line
- In each case you understand what it is that you are choosing between, but you don't know the outcome of that choice
 - In fact, many things can happen, you just don't know which one

- We are going to differentiate between two different ways in which the future may not be known
 - Horse races
 - Roulette wheels
- What is the difference?

- When playing a roulette wheel the probabilities are **known**
 - Everyone agrees on the likelihood of black
 - So we (the researcher) can treat this as something we can observe
 - Probabilities are objective
 - This is a situation of **risk**

- When betting on a horse race the probabilities are **unknown**
 - Different people may apply different probabilities to a horse winning
 - We cannot directly observe a person's beliefs
 - Probabilities are subjective
 - This is a situation of **uncertainty (or ambiguity)**

- We will focus today on choice under risk
- Let's begin by formally defining the objects of choice
- Let X be a finite prize space with N elements
- $\Delta(X)$ the set of probability measures on X

Definition

Let X be some finite prize space, The set $\Delta(X)$ of lotteries on X is the set of all functions $p : X \rightarrow [0, 1]$ such that

$$\sum_{x \in X} p(x) = 1$$

- We will consider preferences over $\Delta(X)$

- What is the cardinality of $\Delta(X)$?
- We will often want to talk about *mixtures* of lotteries

$$\begin{aligned}r &= \alpha p + (1 - \alpha)q \\ \Rightarrow r(x) &= \alpha p(x) + (1 - \alpha)q(x)\end{aligned}$$

- In fact, many of the results that we prove will be special cases of a mathematical result called the **mixture space theorem**
- We will use δ_x to mean the degenerate lottery on prize x
- Sometimes we will abuse notation and use

$$\alpha x + (1 - \alpha)y \text{ to mean } \alpha\delta_x + (1 - \alpha)\delta_y$$

- It is going to be important for our interpretation to make sense that we set up lotteries in the right way
- The only thing that can matter for preferences is the distribution of outcomes
- Consider the following Rubinstein example
 - 50% probability of rain
 - Two prizes $X = \{umbrella, no\ umbrella\}$
- You would not be able to tell my your preferences over a lottery over X *unless you know the 'correlation' between lottery outcome and prizes'*
 - A lottery that gave you an umbrella in the rain and no umbrella otherwise would assign 50% probability to getting an umbrella
 - So would a lottery that gave you no umbrella in the rain and umbrella otherwise
 - But you would not be indifferent between the two....
 - Need to redefine the prize space.....

- So, how should you make choices under risk?
- Let's consider the following (very boring) fairground game
 - You flip a coin
 - If it comes down heads you get \$10
 - If it comes down tails you get \$0
- What is the maximum amount x that you would pay in order to play this game?

Approach 1: Expected Value

- You have the following two options
 - ① Not play the game and get \$0 for sure
 - ② Play the game and get $-\$x$ with probability 50% and $\$10 - x$ with probability 50%

- Approach 1: Expected value

- The expected amount that you would earn from playing the game is

$$0.5(-x) + 0.5(10 - x)$$

- This is bigger than 0 if

$$\begin{aligned} 0.5(-x) + 0.5(10 - x) &\geq 0 \\ 5 &\geq x \end{aligned}$$

- Should pay at most \$5 to play the game

The St. Petersburg Paradox

- This was basically the accepted approach until Daniel Bernoulli suggested the following modification of the game
 - Flip a coin
 - If it comes down heads you get \$2
 - If tails, flip again
 - If that coin comes down heads you get \$4
 - If tails, flip again
 - If that comes down heads, you get \$8
 - Otherwise flip again
 - and so on
- How much would you pay to play this game?

The St. Petersburg Paradox

- The expected value is

$$\begin{aligned} & \frac{1}{2}\$2 + \frac{1}{4}\$4 + \frac{1}{8}\$8 + \frac{1}{16}\$16 + \dots \\ = & \$1 + \$1 + \$1 + \$1 + \dots \\ = & \infty \end{aligned}$$

- So you should pay an infinite amount of money to play this game
- Which is why this is the St. Petersburg **paradox**

- So what is going wrong here?
- Consider the following example:

Example

Say a pauper finds a magic lottery ticket, that has a 50% chance of \$1 million and a 50% chance of nothing. A rich person offers to buy the ticket off him for \$499,999 for sure. According to our 'expected value' method', the pauper should refuse the rich person's offer!

The St. Petersburg Paradox

- It seems ridiculous (and irrational) that the pauper would reject the offer
- Why?
- Because the difference in life outcomes between \$0 and \$499,999 is massive
 - Get to eat, buy clothes, etc
- Whereas the difference between \$499,999 and \$1,000,000 is relatively small
 - A third pair of silk pyjamas
- Thus, by keeping the lottery, the pauper risks losing an awful lot (\$0 vs \$499,999) against gaining relatively little (\$499,999 vs \$1,000,000)

- Bernoulli argued that people should be maximizing expected **utility** not expected **value**
 - $u(x)$ is the expected utility of an amount x
- Moreover, marginal utility should be **decreasing**
 - The value of an additional dollar gets lower the more money you have
- For example

$$u(\$0) = 0$$

$$u(\$499,999) = 10$$

$$u(\$1,000,000) = 16$$

- Under this scheme, the pauper should choose the rich person's offer as long as

$$\frac{1}{2}u(\$1,000,000) + \frac{1}{2}u(\$0) < u(\$499,999)$$

- Using the numbers on the previous slide, LHS=8, RHS=10
 - Pauper should accept the rich person's offer
- Bernoulli suggested $u(x) = \ln(x)$
 - Also explains the St. Petersburg paradox
 - Using this utility function, should pay about \$64 to play the game

- Notice also that expected utility is also a more **general** model than expected value maximization
- The latter can be applied only to cases in which the prize space is amounts of money

- Expected Utility Theory is the workhorse model of choice under risk
- Unfortunately, it is another model which has something unobservable
 - The utility of every possible outcome of a lottery
- So we have to figure out how to identify its observable implications

Definition

A preference relation \succeq on lotteries on some finite prize space X have an expected utility representation if there exists a function $u : X \rightarrow \mathbb{R}$ such that

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

- Notice that preferences are on $\Delta(X)$ but utility numbers are on X
 - Sometimes called Bernoulli numbers

- What needs to be true about preferences for us to be able to find an expected utility representation?
- An **expected utility** representation is still a **utility representation**
- So we still need \succsim to be a preference relation - i.e.
 - Complete
 - Transitive

- Unsurprisingly, this is not enough
- We need two further axioms
 - ① The Independence Axiom
 - ② Continuity

The Independence Axiom

Question: Think of two different lotteries, p and q . Just for concreteness, let's say that p is a 25% chance of winning an apple and a 75% chance of winning a banana, while q is a 75% chance of winning an apple and a 25% chance of winning a banana. Say you prefer the lottery p to the lottery q . Now I offer you the following choice between option 1 and 2

- ① I flip a coin. If it comes up heads, then you get p . Otherwise you get the lottery that gives you celery for sure
- ② I flip a coin. If it comes up heads, you get q . Otherwise you get the lottery that gives you celery for sure

Which do you prefer?

The Independence Axiom

- The independence axiom (effectively) says that if you must prefer p to q you must prefer option 1 to option 2
 - If I prefer p to q , I must prefer a mixture of p with another lottery to q with another lottery

The Independence Axiom $p \succeq q$ implies that, for any other lottery r and number $0 < \alpha \leq 1$ then

$$\alpha p + (1 - \alpha)r \succeq \alpha q + (1 - \alpha)r$$

- Note that the analogy I just gave isn't precise because our world doesn't include two stage lotteries of the type on the previous slide
- Yet independence is often used as a normative axiom
 - For example, it can be used as a decision making tool
 - If you agree with it, then I can ask you some simple questions, and tell you how to behave in more complex situations
- So let's try to construct a normative argument
- We will come back to whether it is accurate descriptively later on

- Consider the following choice scenarios
 - ① When choosing between lottery p and q you prefer p
 - ② Say now I will first flip a coin and that with prob. α you get r , and then you have no choice to make. Otherwise, you get to choose between p and q
 - ③ Say now that you have to commit to a choice of p or q before the coin is flipped
 - ④ Finally I ask you to choose between $\alpha p + (1 - \alpha)r$ and $\alpha q + (1 - \alpha)r$
- We want to conclude that (1) implies you prefer $\alpha p + (1 - \alpha)r$ in (4)
 - $1 \Rightarrow 2$ - history independence
 - $2 \Rightarrow 3$ - time consistency
 - $3 \Rightarrow 4$ - reduction of compound lotteries

- Notice that, while the independence axiom may seem intuitive, that is dependent on the setting
 - Maybe you prefer ice cream to gravy, but you don't prefer ice cream mixed with steak to gravy mixed with steak
- What goes wrong with the previous argument?

- Note that what independence is really buying us here is **linearity in the space of lotteries**

$$x \sim y \Rightarrow \alpha x + (1 - \alpha)y \sim x$$

- One thing that this rules out strict preference for randomization - i.e. we cannot have

$$x \sim y \text{ and } \alpha x + (1 - \alpha)y \succ x$$

Can you think of cases in which a strict preference for randomization makes sense?

- Does this mean the utility numbers assigned to prizes have to be linear?
 - No! We can have concavity in that space (or convexity)
 - This is implicit in how we have defined mixing

- The other axiom we need is more technical

The **Continuity Axiom** For all lotteries p , q and r such that $p \succ q \succ r$, there must exist an a and b in $(0, 1)$ such that

$$ap + (1 - a)r \succ q \succ bp + (1 - b)r$$

- Do we like it?
 - Can't be tested
 - Basically means no prize is infinitely good and infinitely bad
 - What if one prize is death, do we still think it is a good idea?
 - Maybe - after all we still cross roads!

The Expected Utility Theorem

- It turns out that these two axioms, when added to the 'standard' ones, are necessary and sufficient for an expected utility representation

Theorem

Let X be a finite set of prizes and $\Delta(X)$ be the set of lotteries on X . Let \succeq be a binary relation on $\Delta(X)$. Then \succeq is complete, transitive and satisfies Independence and Continuity if and only if there exists a $u : X \rightarrow \mathbb{R}$ such that, for any $p, q \in \Delta(X)$,

$$\text{if and only if } \sum_{x \in X} p_x u(x) \geq \sum_{x \in X} q_x u(x)$$

The Expected Utility Theorem

- Proof?
- Necessity you can do yourself
- Sufficiency relies on two key lemmas

Lemma If \succsim is a preference relation that satisfies Independence then $p \succ q$ and $0 \leq \beta < \alpha \leq 1$ implies

$$\alpha p + (1 - \alpha)q \succ \beta p + (1 - \beta)q$$

Lemma If \succsim is a preference relation that satisfies Independence and Continuity then $p \succ q \succsim r$ and $p \succ r$ implies that there exists a unique α^* such that

$$q \sim \alpha^* p + (1 - \alpha^*)r$$

- Remember that when we talked about 'standard' utility theory, the numbers themselves didn't mean very much
- Only the order mattered
- So, for example

$$u(a) = 1 \quad v(a) = 1$$

$$u(b) = 2 \quad v(b) = 4$$

$$u(c) = 3 \quad v(c) = 9$$

$$u(d) = 4 \quad v(c) = 16$$

- Would represent the same preferences

- Is the same true here?
- No!
- According to the first preferences

$$\frac{1}{2}u(a) + \frac{1}{2}u(c) = 2 = u(b)$$

and so

$$\frac{1}{2}a + \frac{1}{2}c \sim b$$

- But according to the second set of utilities

$$\frac{1}{2}v(a) + \frac{1}{2}v(c) = 5 > v(b)$$

and so

$$\frac{1}{2}a + \frac{1}{2}c \succ b$$

- So we have to take utility numbers more seriously here
 - Magnitudes matter
- How much more seriously?

Theorem

Let \succeq be a set of preferences on $\Delta(X)$ and $u : X \rightarrow \mathbb{R}$ form an expected utility representation of \succeq . Then $v : X \rightarrow \mathbb{R}$ also forms an expected utility representation of \succeq if and only if

$$v(x) = au(x) + b \quad \forall x \in X$$

for some $a \in \mathbb{R}_{++}$, $b \in \mathbb{R}$

Proof.

Homework

