Expected Utility Theory

Mark Dean

Behavioral Economics
Spring 2017

- Up until now, we have thought of subjects 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

Risk vs Uncertainty

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

Risk vs Uncertainty

- 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**

Risk vs Uncertainty

- 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)**

Choices Under Risk

- 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

- 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 expected value is

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

- 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

$$\begin{aligned} u(\$0) &= 0 \\ u(\$499,999) &= 10 \\ u(\$1,000,000) &= 16 \end{aligned}$$

- 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 persons offer
- Bernoulli suggested $u(x) = \ln(x)$
  - Also explains the St. Petersberg paradox
  - Using this utility function, should pay about \$64 to play the game

- Notice that if people
  - Maximize expected utility
  - Have decreasing marginal utility (i.e. utility is concave)
- They will be **risk averse**
  - Will always reject a lottery in favor of receiving its expected value for sure

- 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 test it
- We have already gone through this process for the model of 'standard' (i.e. not expected) utility maximization
- Is this enough for expected utility maximization?

Data

- In order to answer this question we need to state what our data is
- We are going to take as our primitve preferences $\succeq$
  - Not choices
  - But we know how to go from choices to preferences, yes?
- But preferences over what?
  - In the beginning we had preferences over 'objects'
  - For temptation and self control we used 'menus'
  - Now 'lotteries'!

Lotteries

- What is a lottery?
- Like any lottery ticket, it gives you a probability of winning a number of prizes
- Let's imagine there are four possible prizes
  - a(pple), b(anana), c(elery), d(ragonfruit)
- Then a lottery is just a probability distribution over those prizes

$$\begin{pmatrix} 0.15 \\ 0.35 \\ 0.5 \\ 0 \end{pmatrix}$$

- This is a lottery that gives 15% chance of winning a, 35% chance of winning b, 50% of winning c and 0% chance of winning d

- More generally, a lottery is any

$$p = \begin{pmatrix} p_a \\ p_b \\ p_c \\ p_d \end{pmatrix}$$

- Such that
  - $p_x \geq 0$
  - $\sum_x p_x = 1$

Expected Utility

- We say that preferences $\succeq$ have an **expected utility representation** if we can
  - Find utilities on **prizes**
  - i.e. $u(a)$, $u(b)$, $u(c)$, $u(d)$
- Such that

$$p \succeq q \text{ if and only if}$$

$$\begin{aligned} & p_a u(a) + p_b u(b) + p_c u(c) + p_d u(d) \\ > \quad & q_a u(a) + q_b u(b) + q_c u(c) + q_d u(d) \end{aligned}$$

- i.e $\sum_x p_x u(x) \geq \sum_x q_x u(x)$

- What needs to be true about preferences for us to be able to find an expected utility representation?
  - Hint: you know a partial answer to this
- An **expected utility** representation is still a **utility representation**
- So preferences must be
  - Complete
  - Transitive
  - Reflexive

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

Question: Think of two different lotteries, p and q. Just for concreteness, let's say that p is a 25% chance of winning the apple and a 75% chance of winning the banana, while q is a 75% chance of winning the apple and a 25% chance of winning the 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 the celery for sure
2. I flip a coin. If it comes up heads, you get q. Otherwise you get the lottery that gives you the celery for sure

Which do you prefer?

The Independence Axiom

- The independence axiom 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 Say a consumer prefers lottery p to lottery q. Then, for any other lottery r and number $0 < \alpha \leq 1$ they must prefer

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

to

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

- Notice that, while the independence axiom may seem intutive, 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

The Archimedean Axiom

- The other axiom we need is more techincal
- It basically says that no lottery is infinitely good or infinitely bad

The Archimedean 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

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

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

$$\begin{array}{rcl} p & \succeq & q \\ \text{if and only if } \sum_{x \in X} p_x u(x) & \geq & \sum_{x \in X} q_x u(x) \end{array}$$

The Expected Utility Theorem

- Proof?
- Do you want us to go through the proof?
- Oh, alright then
- Actually, Necessity is easy
    - You will do it for homework
- Sufficiency is harder
    - Will sketch it here
    - You can ignore for exam purposes

- Step 1
    - Find the best prize - in other words the prize such that getting that prize for sure is preferred to all other lotteries. Give that prize utility 1 (for convenience, let's say that a is the best prize)
- Step 2
    - Find the worst prize - in other words the prize such that all lotteries are preferred to getting that prize for sure. Give that prize utility 0 (for convenience, let's say that d is the worse prize)
- Step 3
    - Show that, if $a > b$, then

$$a\delta_a + (1-a)\delta_d \succ b\delta_a + (1-b)\delta_d$$

where δ_x is the lottery that gives prize x for sure (this is intuitively obvious, but needs to be proved from the independence axiom)

- Step 4
  - For other prizes (e.g. b), find the probability λ such that the consumer is indifferent between getting apples with probability λ and dragonfruit with probability $(1-\lambda)$, and bananas for sure. Let $u(b)=\lambda$. i.e.

$$\begin{pmatrix} 0 \\ 1 \\ 0 \\ 0 \end{pmatrix} \sim u(b) \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix} + (1-u(b)) \begin{pmatrix} 0 \\ 0 \\ 0 \\ 1 \end{pmatrix}$$

    (for us to know such a λ exists requires the Archimedean axiom)

- Step 5
  - Do the same for c, so

$$\begin{pmatrix} 0 \\ 0 \\ 1 \\ 0 \end{pmatrix} \sim u(c) \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix} + (1-u(c)) \begin{pmatrix} 0 \\ 0 \\ 0 \\ 1 \end{pmatrix}$$

- So now we have found utility numbers for every prize
- All we have to do is show that $p \succeq q$ if and only if $\sum_{x \in X} p_x u(x) \geq \sum_{x \in X} q_x u(x)$
- Let's do a simple example

$$p = \begin{pmatrix} 0 \\ 0.25 \\ 0.75 \\ 0 \end{pmatrix}, \quad q = \begin{pmatrix} 0 \\ 0.75 \\ 0.25 \\ 0 \end{pmatrix}$$

- First, notice that

$$p = \begin{pmatrix} 0 \\ 0.25 \\ 0.75 \\ 0 \end{pmatrix} = 0.25 \begin{pmatrix} 0 \\ 1 \\ 0 \\ 0 \end{pmatrix} + 0.75 \begin{pmatrix} 0 \\ 0 \\ 1 \\ 0 \end{pmatrix}$$

- But

- But

$$\begin{pmatrix} 0 \\ 1 \\ 0 \\ 0 \end{pmatrix} \sim u(b) \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix} + (1 - u(b)) \begin{pmatrix} 0 \\ 0 \\ 0 \\ 1 \end{pmatrix}$$

and

$$\begin{pmatrix} 0 \\ 0 \\ 1 \\ 0 \end{pmatrix} \sim u(c) \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix} + (1 - u(c)) \begin{pmatrix} 0 \\ 0 \\ 0 \\ 1 \end{pmatrix}$$

$$
\begin{aligned}
p \;\sim\; & 0.25\left(u(b)\begin{pmatrix}1\\0\\0\\0\end{pmatrix} + (1-u(b))\begin{pmatrix}0\\0\\0\\1\end{pmatrix}\right) \\
& +0.75\left(u(c)\begin{pmatrix}1\\0\\0\\0\end{pmatrix} + (1-u(c))\begin{pmatrix}0\\0\\0\\1\end{pmatrix}\right)
\end{aligned}
$$

$$
\begin{aligned}
= \quad & (0.25u(b) + 0.75u(c)) \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix} + \\
& (1 - 0.25u(b) - 0.75u(c)) \begin{pmatrix} 0 \\ 0 \\ 0 \\ 1 \end{pmatrix}
\end{aligned}
$$

- So p is indifferent to a lottery that puts probability

$$(0.25u(b) + 0.75u(c))$$

  on the best prize (and the remainder on the worst prize)
- **But this is just the expected utility of** p
- Similarly q is indfferent to a lottery that puts

$$(0.75u(b) + 0.25u(c))$$

  on the best prize
- **But this is just the expected utility of** q

- So p will be preferred to q if the expected utility of p is higher than the expected utility of q
- This is because this means that p is indifferent to a lottery which puts a higher weight on the best prize than does q
- QED (ish)

Expected Utility Numbers

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

$$
\begin{aligned}
u(a) &= 1 \quad v(a) = 1 \\
u(b) &= 2 \quad v(b) = 4 \\
u(c) &= 3 \quad v(c) = 9 \\
u(d) &= 4 \quad v(c) = 16
\end{aligned}
$$

- 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 $u : X \to \mathbb{R}$ form an expected utility representation of $\succeq$. Then $v : X \to \mathbb{R}$ also forms an expected utility representation of $\succeq$ if and only if

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

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

Proof.

Homework □