Expected Utility Theory

Mark Dean

GR5211 - Microeconomic Analysis 1

- 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

- 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)**

- 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)

Marginal Utility

- 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

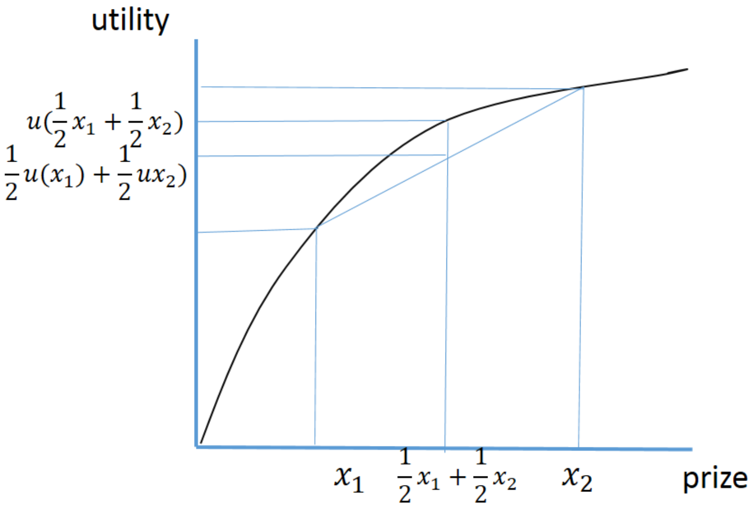
utility
$u(\frac{1}{2}x_1 + \frac{1}{2}x_2)$
$\frac{1}{2}u(x_1) + \frac{1}{2}ux_2)$
x_1
$\frac{1}{2}x_1 + \frac{1}{2}x_2$
x_2
prize

- 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'!

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

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

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

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 $\sum_x p_x u(x) \geq \sum_x q_x u(x)$

Definition

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

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

- Notice that preferences are on $\Delta(X)$ but utility numbers are on 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 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 $p, q \in \Delta(X)$,

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

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

Proof

- Key to the proof is the following lemma

Lemma If $\succeq$ is complete, reflexive, transitive and satisfies the Independence and Archimedean axioms then

1. $p \succ q$ and $0 \leq \alpha < \beta \leq 1$ implies

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

2. $p \succeq q \succeq r$ and $p \succ r$ implies that there exists a unique α^* such that

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

- 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

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

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

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

The Expected Utility Theorem

- 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**

The Expected Utility Theorem

- 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 $u : X \rightarrow \mathbb{R}$ 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

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

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

Proof.

Homework □

Risk Aversion

- We motivated the move from expected value maximization to expected **utility** maximization on the basis of risk aversion
- Does EU imply risk aversion?
- No!
- Consider someone who has $u(x) = x$
  - They will be risk neutral
- Consider someone who has $u(x) = x^2$
  - They will be risk loving
- So risk attitude has something to do with the shape of the utility function

- For this section we will think about lotteries with monetary prizes
- Let δ_x be the lottery that gives prize x for sure and $E(p)$ be the expected value of a lottery p

Definition

We say that a decision maker is risk averse if, for every lottery p

$$\delta_{E(p)} \succeq p$$

We say they are risk neutral if

$$\delta_{E(p)} \sim p$$

We say they are risk loving if

$$\delta_{E(p)} \preceq p$$

- We can say the same thing a different way

Definition

The **certainty equivalence** of a lottery p is the amount c such that

$$\delta_c \sim p$$

The risk premium is

$$E(p) - c$$

Lemma

For a decision maker whose preferences are strictly monotonic in money

1. *They are risk averse if and only if for any p the risk premium is weakly positive*
2. *They are risk neurtal if and only if for any p the risk premium is zero*
3. *They are risk loving if and only if for any p the risk premium is weakly negative*

Risk Aversion and Utility Curvature

- We have made the claim that there is a link between risk aversion and the curvature of the utility function

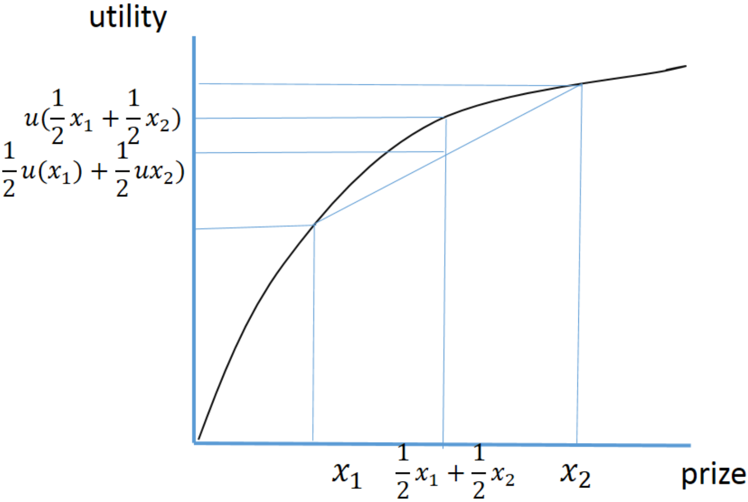

Risk Aversion and Utility Curvature

- We can make this statement tight

Theorem

An expected utility maximizer

1. *Is risk averse if and only if u is concave*
2. *Is risk neutral if and only if u is linear*
3. *Is risk loving if and only if u is convex*

Proof.

Comes straight from Jensen's inequality: for a random variable x and a concave function u

$$E(u(x)) \leq u(E(x))$$

□

Measuring Risk Aversion

- We might want a way of measuring risk aversion from the utility function
- Intuitively, the more 'curvy' the utility function, the more risk averse
- How do we measure curvature?
- The second derivative $u''(x)$!
- Is this a good measure?
- No, because we can change the utility function in such a way that we don't change the underlying preferences, and change $u''(x)$

- One way round this problem is to use the **Arrow-Pratt** measure of **absolute** risk aversion

$$A(x) = \frac{-u''(x)}{u'(x)}$$

- This measure has some nice properties
  1. If two utility functions represent the same preferences then they have the same A for every x
  2. It measures risk aversion in the sense that the following two statements are equivalent
     - The utility function u has a higher Arrow Pratt measure than utility function v for every x
     - Utility function u gives a higher risk premium than utility function v for every p

- Why is it called a measure of **absolute** risk aversion?
- To see this, let's think of a function for which $A(x)$ is constant

$$u(x) = 1 - e^{-ax}$$

  - Note $u'(x) = ae^{-ax}$ and $u''(x) = -a^2 e^{-ax}$ so $A(x) = a$
- This is a constant absolute risk aversion (CARA) utility function

- Claim: for CARA utility functions, adding a constant amount to each lottery doesn't change risk attitues
- i.e if $\delta_x \succeq p$ then δ_{x+z} is preferred to a lottery p' which adds an amount z to each prize in p
- To see this note that

$$
\begin{aligned}
u(x) &\geq \sum_y p(y)u(y) \\
1 - e^{-ax} &\geq \sum_y p(y)\left(1 - e^{-ay}\right) \\
&\Rightarrow 1 - e^{-ax} \geq 1 - \sum_y p(y)e^{-ay} \\
e^{-az} - e^{-ax}e^{-az} &\geq e^{-az} - \sum_y p(y)e^{-ay}e^{-az} \\
&\Rightarrow 1 - e^{-a(x+z)} \geq \sum_y p(y)\left(1 - e^{-a(y+z)}\right) \\
&\Rightarrow u(x+z) \geq \sum_y p(y)u(y+z)
\end{aligned}
$$

- Is this a sensible property?
- Maybe not
- Means that you should have the same attitude to a gamble between winning \$100 or losing \$75 whether you are a student earning \$20,000 a year or a professor earning millions!
- Perhaps a more useful measure is **relative** risk aversion

$$R(x) = xA(x) = -\frac{xu''(x)}{u'(x)}$$

Relative Risk Aversion

- An example of a Constant Relative Risk Aversion measure is

$$u(x) = \frac{x^{1-\rho} - 1}{1 - \rho}$$

- Note that $u'(x) = x^{-\rho}$, $u''(x) = -\rho x^{-\rho-1}$ and so $R(x) = \rho$
- CRRA utility functions have the property that proportional changes in prizes don't affect risk attitudes
- i.e if $\delta_x \succeq p$ then $\delta_{\alpha x}$ is preferred to a lottery p' which multiplies each prize in p by $\alpha > 0$

- To see this note that

$$
\begin{aligned}
u(x) &\geq \sum_y p(y)u(y) \\
&\Rightarrow \frac{x^{1-\rho}-1}{1-\rho} \geq \frac{\sum_y p(y)y^{1-\rho}-1}{1-\rho} \\
&\Rightarrow x^{1-\rho} \geq \sum_y p(y)y^{1-\rho} \\
&\Rightarrow \alpha^{1-\rho}x^{1-\rho} \geq \sum_y p(y)\alpha^{1-\rho}y^{1-\rho} \\
&\Rightarrow \frac{(\alpha x)^{1-\rho}-1}{1-\rho} \geq \frac{\sum_y p(y)\left(\alpha y\right)^{1-\rho}-1}{1-\rho} \\
u(\alpha x) &\geq \sum_y p'(y)u(y)
\end{aligned}
$$

Are People Expected Utility Maximizers?

- Because of the work we have done above, we know what the 'behavioral signature' is of EU
  - The independence axiom
- Essentially this is picking up on the fact that EU demands preferences to be linear in probabilities
- Does this hold in experimental data?

The Common Ratio Effect

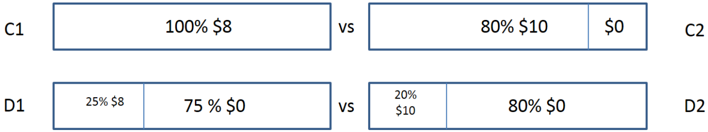

- What would you choose?
- Many people choose C1 and D2

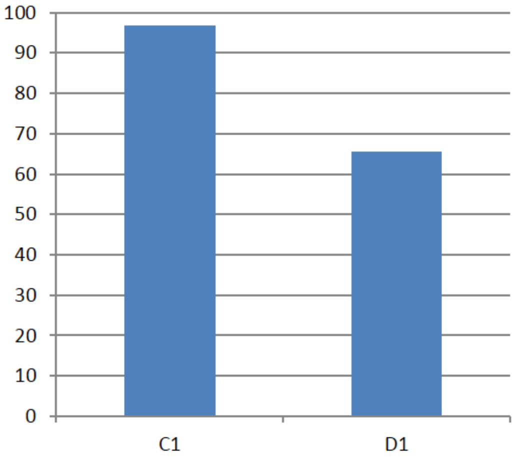
100
90
80
70
60
50
40
30
20
10
0
C1
D1

- This is a violation of the independence axiom
- Why?
- Because

$$\begin{aligned} D1 &= 0.25C1 + 0.75R \\ D2 &= 0.25C2 + 0.75R \end{aligned}$$

  where R is the lottery which pays 0 for sure
- Thus independence means that

$$C1 \succeq C2 \Rightarrow D1 \succeq D2$$

The Common Consequence Effect

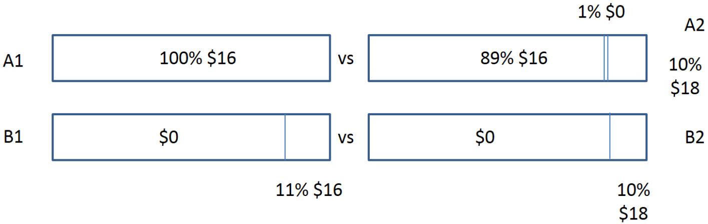

- What would you choose?
- Many people choose A1 and B2

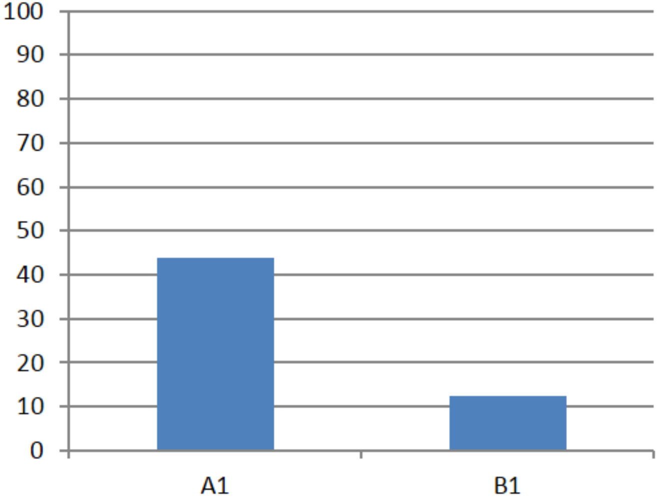
100
90
80
70
60
50
40
30
20
10
0
A1
B1

Explanations

- What do you think is going on?
- Many alternative models have been proposed in the literature
  - Disappointment: Gul, Faruk, 1991. "A Theory of Disappointment Aversion,"
  - Salience: Pedro Bordalo & Nicola Gennaioli & Andrei Shleifer, 2012. "Salience Theory of Choice Under Risk,"
- One of the most widespread and straightforward is **probability weighting**

Probability Weighting

- Maybe the problem that the Allais paradox highlights is that people do not 'believe' the probabilities that are told to them
  - For example they treat a 1% probability of winning $0 as if it is more likely than that
    - ‘I am unlucky, so the bad outcome is more likely to happen to me’
  - The difference between 0% and 1% seems bigger than the difference between 89% and 90%
- This is the idea behind the probability weighting model.

Simple Probability Weighting Model

- Approach 1: Simple probability weighting
- Let's start with expected utility

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

- And allow for probability weighting

$$V(p) = \sum_{x \in X} \pi(p(x))u(x)$$

  Where π is the probability weighting function
- This can explain the Allais paradox
  - For example if $\pi(0.01) = 0.05$

Simple Probability Weighting Model

- However, the simple probability weighting model is not popular
- For two reasons
  1. It leads to violations of stochastic dominance
  2. It doesn't really capture the idea of 'pessimism'

Pessimism

- Think back to the Allais paradox

$$\begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix} \succ \begin{pmatrix} 0.01 \\ 0.89 \\ 0.1 \end{pmatrix}$$

- It seems as if the 1% probability of \$0 is being overweighted
- Is this just because it is a 1% probability?
- Or is it because it is a 1% probability **of the worst prize**
- If it is the latter, this is something that the simple probability weighting model cannot capture
  - Weights are only based on probability

- Consider the following two examples

Example

Lottery p :49% chance of \$10, 49% of winning \$0, 2% chance of winning \$5

Example

Lottery p :49% chance of \$10, 49% of winning \$0, 2% chance of losing \$1000

- Would you 'weigh' the 2% probability the same in each case?
  - Arguably not
  - If you were pessimistic then you might think that 2% is 'more likely' in the latter case than in the former
  - Can't be captured by the simple probability weighting model

- Because of these two concerns, the simple probability weighting model is rarely used
- Instead people tend to use **rank dependent utility** (sometimes also called cumulative probability weighting)
- Probability weighting depends on
  - The **probability** of a prize
  - Its **rank** in the lottery - i.e. how many prizes are better or worse than it
- In practice this is done by applying weights **cumulatively**
- Here comes the definition
  - It looks scary, but don't panic!

Definition

A decision maker's preferences $\succeq$ over $\Delta(X)$ can be represented by a rank dependant utility model if there exists a utility function $u : X \rightarrow \mathbb{R}$ and a cumulative probability weighting function $\psi : [0, 1] \rightarrow [0, 1]$ such that $\psi(0) = 0$ and $\psi(1) = 1$, such that the function $U : \Delta(X) \rightarrow \mathbb{R}$ represents $\succeq$, where $U(p)$ is constructed in the following way:

1. The prizes of p are ranked $x_1, x_2, \ldots, x_n$ such that $x_1 \succ x_2 \cdots \succ x_n$
2. $U(p)$ is determined as

$$U(p) = \psi(p_1)u(x_1) + \sum_{i=2}^{n} \left(\psi\left(\sum_{j=1}^{i} p_j \right) - \psi\left(\sum_{k=1}^{i-1} p_k \right) \right) u(x_i)$$

- Let's go through an example: for prizes $10 > 5 > 0$ let p be equal to

$$\begin{pmatrix} 0.1 \\ 0.7 \\ 0.2 \end{pmatrix}$$

- How do we apply RDU?

- Well, first note that there are three prizes, so we can rewrite the expression above as

$$\begin{aligned} U(p) &= \psi(p_1)u(x_1) \\ &+ (\psi(p_1+p_2) - \psi(p_1))\, u(x_2) \\ &+ (\psi(p_1+p_2+p_3) - \psi(p_1+p_2))\, u(x_3) \end{aligned}$$

- The weight attached to the best prize is the weight of p_1
- The weight attached to the second best prize is the weight on the probability of
  - Getting something at least as good as the second prize
  - Minus the probability of getting something better than the second prize
  - And so on
- Notice that if ψ is the identity function this is just expected utility

- In this specific case

$$\begin{aligned} U(p) &= \psi(p_1)u(x_1) \\ &+ \left(\psi\left(p_1+p_2\right)-\psi\left(p_1\right)\right)u(x_2) \\ &+ \left(\psi\left(p_1+p_2+p_3\right)-\psi\left(p_1+p_2\right)\right)u(x_3) \end{aligned}$$

- Becomes

$$\begin{aligned} U(p) &= \psi(0.1)u(10) \\ &+ \left(\psi\left(0.8\right)-\psi\left(0.1\right)\right)u(5) \\ &+ \left(\psi\left(1\right)-\psi\left(0.8\right)\right)u(0) \end{aligned}$$

- In the first class we drew a distinction betweem
  - Circumstances of **Risk** (roulette wheels)
  - Circumstances of **Uncertainty** (horse races)
- So far we have been talking about roulette wheels
- Now horse races!

Risk vs Uncertainty

- Remember the key difference between the two
- Risk: Probabilities are **observable**
  - There are 38 slots on a roulette wheel
  - Someone who places a $10 bet on number 7 has a lottery with pays out $350 with probability 1/38 and zero otherwise
  - (Yes, this is not a fair bet)
- Uncertainty: Probabilities are **not observable**
  - Say there are 3 horses in a race
  - Someone who places a $10 bet on horse A does not necessarily have a 1/3 chance of winning
  - Maybe their horse only has three legs?

- If we want to model situations of uncertainty, we cannot think about preferences over **lotteries**
- Because we don't know the probabilities
- We need a different set up
- We are going to thing about **acts**
- What is an act?

States of the World

- First we need to define **states of the world**
- We will do this with an example
- Consider a race between three horses
  - A(rchibald)
  - B(yron)
  - C(umberbach)
- What are the possible oucomes of this race?
  - Excluding ties

States of the World

State	Ordering
1	A, B ,C
2	A, C, B
3	B, A, C
4	B, C, A
5	C, A, B
6	C, B, A

Acts

- This is what we mean by the states of the world
  - An exclusive and exhaustive list of all the possible outcomes in a scenario
- An **act** is then an action which is defined by the oucome it gives in each state of the world
- Here are two examples
  - Act f: A \$10 even money bet that Archibald will win
  - Act g: A \$10 bet at odds of 2 to 1 that Cumberbach will win

State	Ordering	Payoff Act f	Payoff Act g
1	A, B ,C	$10	-$10
2	A, C, B	$10	-$10
3	B, A, C	-$10	-$10
4	B, C, A	-$10	-$10
5	C, A, B	-$10	$20
6	C, B, A	-$10	$20

- So, how would you choose between acts f and g?
- SEU assumes the following:

1. Figure out the probability you would associate with each state of the world
2. Figure out the utility you would gain from each prize
3. Figure out the expected utility of each act according to those probabilities and utilities
4. Choose the act with the highest utility

- So, in the above example
- Utility from f :

$$\begin{aligned}&[\pi(ABC)+\pi(ACB)]\,u(10)\\&+[\pi(BAC)+\pi(BCA)]\,u(-10)\\&+[\pi(CBA)+\pi(CAB)]\,u(-10)\end{aligned}$$

where π is the probability of each act

- Utility from g :

$$\begin{aligned}&[\pi(ABC)+\pi(ACB)]\,u(-10)\\&+[\pi(BAC)+\pi(BCA)]\,u(-10)\\&+[\pi(CBA)+\pi(CAB)]\,u(20)\end{aligned}$$

- Assuming utility is linear f is preferred to g if

$$\frac{[\pi(ABC) + \pi(ACB)]}{[\pi(CBA) + \pi(CAB)]} \geq \frac{3}{2}$$

- Or the probability of A winning is more than $3/2$ times the probability of C winning

Definition

Let X be a set of prizes, Ω be a (finite) set of states of the world and F be the resulting set of acts (i.e. F is the set of all functions $f : \Omega \to X$). We say that preferences $\succeq$ on the set of acts F has a subjective expected utility representation if there exists a utility function $u : X \to \mathbb{R}$ and probability function $\pi : \Omega \to [0, 1]$ such that $\sum_{\omega \in \Omega} \pi(\omega) = 1$ and

$$
\begin{aligned}
f &\succeq g \\
&\Leftrightarrow \sum_{\omega \in \Omega} \pi(\omega) u\left(f(\omega)\right) \geq \sum_{\omega \in \Omega} \pi(\omega) u\left(g(\omega)\right)
\end{aligned}
$$

- Notes
  - Notice that we now have **two** things to recover: Utility and preferences
  - Axioms beyond the scope of this course: has been done twice - first by Savage[1] and later (using a trick to make the process a lot simpler) by Anscombe and Aumann[2]
  - Utility pinned down to positive affine transform
  - Probabilities are unique

[1] Savage, Leonard J. 1954. The Foundations of Statistics. New York, Wiley.

[2] Anscombe, F. J.; Aumann, R. J. A Definition of Subjective Probability. The Annals of Mathematical Statistics 34 (1963), no. 1, .

- Unfortunately, while simple and intuitive, SEU theory has some problems when it comes to describing behavior
- These problems are most elegantly demostrated by the Ellsberg paradox
  - This thought experiment has sparked a whole field of decision theory

The Ellsberg Paradox - A Reminder

- Choice 1: The 'risky bag'
  - Fill a bag with 20 red and 20 black tokens
  - Offer your subject the opportunity to place a \$10 bet on the color of their choice
  - Then elicit the amount x such that the subject is indifferent between playing the gamble and receiving \$$x$ for sure.
- Choice 2: The 'ambiguous bag'
  - Repeat the above experiment, but provide the subject with no information about the number of red and black tokens
  - Then elicit the amount y such that the subject is indifferent between playing the gamble and receiving \$$y$ for sure.

- Typical finding
  - $x >> y$
  - People much prefer to bet on the risky bag
- This behavior cannot be explained by SEU?
- Why?

- What is the utility of betting on the risky bag?
- The probability of drawing a red ball is the same as the probability of drawing a black ball at 0.5
- So whichever act you choose to bet on, the utility of the gamble is

$$0.5u(\$10)$$

- What is the utility of betting on the ambiguous bag?
- Here we need to apply SEU
- What are the states of the world?
  - Red ball is drawn or black ball is drawn
- What are the acts?
  - Bet on red or bet on black

State	r	b
red	10	0
black	0	10

- How do we calculate the utility of these two acts?
  - Need to decide how likely each state is
  - Assign probabilities $\pi(r) = 1 - \pi(b)$
  - Note that these do **not** have to be 50%
  - Maybe you think I like red chips!

- Utility of betting on the red outcome is therefore

$$\pi(r)u(\$10)$$

- Utility of betting on the black outcome is

$$\pi(b)u(\$10) = (1 - \pi(r))u(\$10)$$

- Because you get to choose which color to bet on, the gamble on the ambiguous urn is

$$\max\{\pi(r)u(\$10), (1 - \pi(r))u(\$10)\}$$

  - is equal to $0.5u(\$10)$ if $\pi(r) = 0.5$
  - otherwise is **greater** than $0.5u(\$10)$
  - should always (weakly) prefer to bet on the ambiguous urn
  - intuition: if you can choose what to bet on, 0.5 is the worst probability

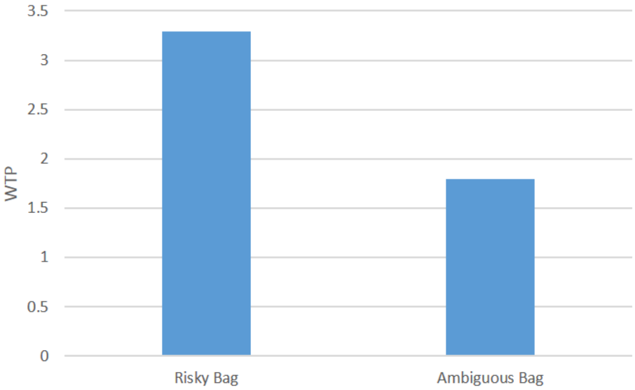

- 61% of my last class exhibited the Ellsberg paradox
- For more details see *Halevy, Yoram. "Ellsberg revisited: An experimental study." Econometrica 75.2 (2007): 503-536.*

- So, as usual, we are left needing a new model to explain behavior
- There have been many such attempts since the Ellsberg paradox was first described
- We will focus on 'Maxmin Expected Utility' by Gilboa and Schmeidler[3]

[3] Gilboa, Itzhak & Schmeidler, David, 1989. "Maxmin expected utility with non-unique prior," Journal of Mathematical Economics, Elsevier, vol. 18(2), pages 141-153, April.

Maxmin Expected Utility

- Maxmin expected utility has a very natural interpretation....
- The world is out to get you!
  - Imagine that in the Ellsberg experiment was run by an evil and sneaky experimenter
  - After you have chosen whether to bet on red or black, they will increase your chances of losing
  - They will sneak some chips into the bag of the **opposite** color to the one you bet on
  - So if you bet on red they will put black chips in and visa versa

- How should we think about this?
- Rather than their being a single probability distribution, there is a **range** of possible distributions
- After you chose your act, you evaluate it using the **worst** of these distributions
- This is maxmin expected utility
  - you **maximize** the **minimum** utility that you can get across different probability distributions
- Has links to robust control theory in engineering
  - This is basically how you design aircraft

Definition

Let X be a set of prizes, Ω be a (finite) set of states of the world and F be the resulting set of acts (i.e. F is the set of all functions $f : \Omega \to X$). We say that preferences $\succeq$ on the set of acts F has a Maxmin expected utility representation if there exists a utility function $u : X \to \mathbb{R}$ and convex set of probability functions Π and

$$\begin{aligned} f &\succeq g \\ &\Leftrightarrow \min_{\pi \in \Pi} \sum_{\omega \in \Omega} \pi(\omega) u\left(f(\omega)\right) \geq \min_{\pi \in \Pi} \sum_{\omega \in \Omega} \pi(\omega) u\left(g(\omega)\right) \end{aligned}$$

Maxmin Expected Utility

- Maxmin expected utility can explain the Ellsberg paradox
  - Assume that $u(x) = x$
  - Assume that you think $\pi(r)$ is between 0.25 and 0.75
  - Utility of betting on the risky bag is $0.5u(x) = 5$
  - What is the utility of betting on red from the ambiguous bag?

$$\min_{\pi(r) \in [0.25, 0.75]} \pi(r) u(\$10) = 0.25u(\$10) = 2.5$$

  - Similary, the utility from betting on black is

$$\min_{\pi(r) \in [0.25, 0.75]} (1 - \pi(r))\, u(\$10) = 0.25u(\$10) = 2.5$$

  - Maximal utility from betting on the ambiguous bag is lower than that from the risky bag