

Risk and Uncertainty 2

Mark Dean

GR6211 - Microeconomic Analysis 1

- We motivated EU theory by appealing to 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 neutral 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



- 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

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



Risk Aversion and Expected Utility

- Note that our definition of risk aversion seems rather weak
 - Only talk about the comparison between risky lotteries and sure things
- We might want our definition to cover other situations
 - e.g. we might want a risk aversion decision maker to dislike mean preserving spreads
- The above result show that we get this 'for free' if we apply our definition of risk aversion to EU maximizers
 - If the DM is an EU maximizer then

$$\begin{array}{l} \text{Risk aversion} \quad \Leftrightarrow \\ \text{Concave Utility} \quad \iff \\ p \quad \succeq \quad q \text{ for all } q \text{ mean preserving spread of } p \end{array}$$

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

The Arrow Pratt Measure

- Claim: for CARA utility functions, adding a constant amount to each lottery doesn't change risk attitudes
- 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

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

- 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

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

- 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

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

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

The Common Ratio Effect



The Common Ratio Effect

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

$$D1 = 0.25C1 + 0.75R$$


$$D2 = 0.25C2 + 0.75R$$

where R is the lottery which pays 0 for sure

- Thus independence means that


$$C1 \succ C2 \Rightarrow D1 \succ D2$$

The Common Consequence Effect



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

The Common Consequence Effect



- 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,"
 - Saliency: Pedro Bordalo & Nicola Gennaioli & Andrei Shleifer, 2012. "Saliency Theory of Choice Under Risk,"
- One of the most widespread and straightforward is **cumulative probability weighting**
 - e.g. Quiggin, J. 1982, 'A theory of anticipated utility'
- Beyond the scope of the course but basic idea
 - Rank prizes from best to worst
 - Reweight the probability of each prize based on its position in the CDF

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

- 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 think about **acts**
- What is an act?

- 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 outcomes of this race?
 - Excluding ties

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

- 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 outcome 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

Subjective Expected Utility Theory

- 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

Subjective Expected Utility Theory

- 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 :

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

Subjective Expected Utility Theory

- 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 $f : \Omega \rightarrow X$). We say that preferences \succeq on the set of acts F has a subjective expected utility representation if there exists a utility function $u : X \rightarrow \mathbb{R}$ and probability function $\pi : \Omega \rightarrow [0, 1]$ such that $\sum_{\omega \in \Omega} \pi(\omega) = 1$ and

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

Notice that we now have **two** things to recover: Utility and preferences

Subjective Expected Utility Theory

- As with expected utility theory, we would like a set of axioms on \succeq which are necessary and sufficient for an SEU representation
- Savage did this in 1954
- In this setting it is pretty tricky
- Why?
 - Because really we still want to impose linearity in probabilities
 - Would like to impose something like the independence axiom
 - But we don't know how to mix acts together

The Anscombe Aumann Set Up

- Anscombe and Aumann came up with a nice trick to get round this in 1963
- Rather than have acts pay off **prizes** in each state, have them pay off **lotteries**
- Let
 - X be a (finite) prize space
 - Ω be a finite state space

Definition

An Anscombe-Aumann act is a function $h : \Omega \rightarrow \Delta(X)$. Let H define the set of all such acts

- In this world we can define mixing

Definition

The mixture of two acts $h, g \in H$, $\alpha h + (1 - \alpha)g \in H$ is the act such that, for each state $\omega \in \Omega$

$$(\alpha h + (1 - \alpha)g)(\omega) = \alpha h(\omega) + (1 - \alpha)g(\omega)$$

The Anscombe Aumann Set Up

- And so define equivalents of the independence and continuity axioms

The Independence Axiom $h \succeq g$ implies that, for any other act f and number $0 < \alpha \leq 1$ then

$$\alpha h + (1 - \alpha)f \succeq \alpha g + (1 - \alpha)f$$

The Continuity Axiom For all acts h, g and f such that $h \succ g \succ f$, there must exist an a and b in $(0, 1)$ such that

$$ah + (1 - a)f \succ g \succ bh + (1 - b)f$$

The Anscombe Aumann Set Up

- If we add to this three standard axioms
 - \succsim is a preference relation
 - \succsim non degenerate
 - Monotonicity: if $h(\omega) \succeq g(\omega) \forall \omega \in \Omega$ then $h \succeq g$
- Then we get a modified SEU representation

Definition

A preference relation \succsim on H has an SEU representation if there exists a utility function $u : X \rightarrow \mathbb{R}$ and probability function $\pi : \Omega \rightarrow [0, 1]$ such that

$$h \succsim g$$

$$\Leftrightarrow \sum_{\omega \in \Omega} \pi(\omega) \left[\sum_{x \in X} h(\omega)(x) u(x) \right] \geq \sum_{\omega \in \Omega} \pi(\omega) \left[\sum_{x \in X} g(\omega)(x) u(x) \right]$$

- Unfortunately, while simple and intuitive, SEU theory has some problems when it comes to describing behavior
- These problems are most elegantly demonstrated 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
 - $x \gg y$
 - 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

The Ellsberg Paradox



- 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 Schmeidler¹

¹Gilboa, 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 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 $f : \Omega \rightarrow X$). We say that preferences \succeq on the set of acts F has a Maxmin expected utility representation if there exists a utility function $u : X \rightarrow \mathbb{R}$ and convex set of probability functions Π and

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

- 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?

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

- Similarly, 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