8 Expected Utility

Thus far, we've characterize preferences over deterministic bundles that an individual consumes with probability one. Now we extend our model of utility maximization to characterize behavior when payoffs are uncertain – because they depend on the actions of others or factors beyond our control.

8.1 The St. Petersburg Paradox

- Consider the following gamble: you flip a coin over and over until the first time it lands on heads. If it (first) lands on heads on the n^{th} flip, you win 2^n dollars.
  - How much would you pay to play this gamble? \$1? \$10? \$100? \$1,000?
- We might expect willingness to pay to be related to the expected payout from a gamble (i.e. the expected value of the lottery). What is the expected value of this particular gamble?

Sequence	**Probability**	**Payout**
H	1/2	2
TH	1/4	$2^2 = 4$
TTH	1/8	$2^3 = 8$
TTTH	1/16	$2^4 = 16$
TTTTH	1/32	$2^5 = 32$
...	...	...
n	$1/2^n$	2^n

- The expected value of the gamble is the sum of the possibly payoffs weighted by their respective probabilities:

$$
\begin{aligned}
EV &= \frac{1}{2}(2) + \frac{1}{4}(4) + \frac{1}{8}(8) + \dots \\
EV &= \frac{1}{2^1}(2^1) + \frac{1}{2^2}(2^2) + \frac{1}{2^3}(2^3) + \dots \\
&= \sum_{n=1}^{\infty} \frac{1}{2^n}\left(2^n\right) \qquad (1) \\
&= \sum_{n=1}^{\infty} 1 \\
&= \infty
\end{aligned}
$$

- In spite of the fact that this is an infinite expected value lottery, very few people are willing to pay a lot for this particular lottery – a fact first noted by the mathematician Nicolaus Bernoulli in 1738.[1]

- Daniel Bernoulli's proposed solution to the paradox: people do not care about the *amount* of money they receive, but instead the *natural log of the amount* of money they receive.

- Bernoulli was the first to point out that the utility/happiness associated with receiving x may concave and not be linear in x: the difference in happiness that results from receiving one thousand dollars instead of zero dollars is larger than the difference in happiness that results from receiving receiving 17 thousand dollars instead of 16 thousand dollars

- The insight that the utility of money may not be linear is at the core of economist's models of attitudes toward risk

[1] Nicolaus was one of several mathematicians in the Bernoulli family. A solution to Nicolaus' paradox was proposed by his cousin Daniel Bernoulli. Nicolaus and Daniel were both nephews of Jabob Bernoulli, for whom the Bernoulli distribution is named. Jacob's brother (and Daniel's father) Johann Bernoulli was a mathematician who married a member of the Curie family, and also went on to train Leonhard Euler.

8.2 Lotteries

- A **lottery** is a set of probabilities that sum to one, each with an associated payoff:

$$L = \{\pi_1, \pi_2, \ldots, \pi_K; x_1, x_2, \ldots, x_K\} \text{ where } \sum_{k=1}^{K} \pi_k = 1 \tag{2}$$

  - So, the outcome of a lottery is a random variable
  - Technically this is a money lottery, and the prizes in a simple lottery could be outcomes other than money. In this course, we are confining our analysis of preferences over risky outcomes to the case where all outcomes are money.
  - We will also use the terms **gamble** and **prospect** to refer to lotteries.
  - It is important to note that the probabilities are known – we can also model situations where decision makers are uncertain of the probabilities, but that is not what we are doing here.

- Examples:

$$L = \{0.5, 0.5; 0, 100\}$$

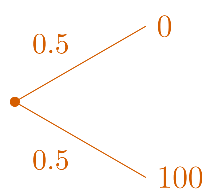

$$L = \{0.25, 0.25, 0.5; 20, 60, 100\}$$

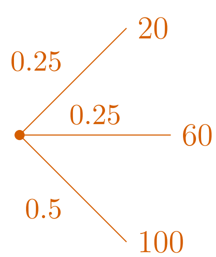

- The **expected value** of a lottery is $\sum_{k=1}^{K} \pi_k x_k$
  - When all payoffs equally likely, $EV(L)$ is average payoff
  - Examples:
    1. $L = \{0.5, 0.5; 0, 100\} \Rightarrow EV(L) = 0.5 \cdot 0 + 0.5 \cdot 100 = 50$
    2. $L = \{0.25, 0.25, 0.5; 20, 60, 100\} \Rightarrow EV(L) = 5 + 15 + 50 = 70$
- A **degenerate lottery** is a lottery with only one possible outcome (which occurs with probability one)

8.3 Expected Utility

- Thus far, our discussions of utility functions have focused on cases where the arguments are deterministic payoffs. When we want to comparisons between risky prospects (i.e. lotteries), we need to impose additional assumptions to guarantee that tastes can be represented by a utility function.
- A **von Neumann-Morgenstern (VNM) exected utility function** is a utility function over lotteries that takes the form:

$$EU(L) = \pi_1 u(x_1) + \pi_2 u(x_2) + \ldots + \pi_K u(x_K) \tag{3}$$

  given a **Bernoulli utility function** $u(x)$ defined over money outcomes
  - A **Bernoulli utility function** is and is not a utility function in the "first half of ECON 251" sense
  - It assigns a utility level $u(x)$ to each monetary payoff x, so in that sense it **is** a utility function
  - However, unlike the utility functions that we've studied before, $u(x)$ and a concave transformation of $u(x)$ do **not** represent the same preferences over risky monetary payoffs
  - The reason goes back to Bernoulli: how we feel about risk depends on the curvature of the Bernoulli utility function; so we can only do linear transformations of Bernoulli utility functions

- If we were only trying to rank degenerate lotteries, then we could still do monotone transformations of Bernoulli utility functions... but we also want to rank lotteries
- Whether one prefers a 50 percent chance of 100 dollars or an 80 percent chance of 60 dollars depends on how much happier one would be with 100 rather than 60 dollars

◦ Expected utility means that utility is "linear in probability"

- A **linear combination** is constructed by multiplying each of a set of terms by a (potentially different) constant then adding up the terms
- In the expected utility framework, the utility of a lottery is a linear combination of the Bernoulli utilities of the possible payoffs/outcomes

◦ We can also use expected utility to model decisions over risky prospects involving non-monetary outcomes by specifying the set of outcomes (prizes) and the consumption utility associated with each prize (sometimes referred to as the Bernoulli utility or the Bernoulli value of the prize)

◦ Example: $u(x) = \sqrt{x}$

Example 1:

$$L = \{1; 25\} \longleftarrow \text{"degenerate lottery"}$$

$$EU(L) = \sqrt{25} = 5$$

Example 2:

$$L = \{0.5, 0.5; 0, 100\}$$

$$EU(L) = 0.5\sqrt{0} + 0.5\sqrt{100} = 5$$

8.4 Risk Aversion

◦ Suppose $u(x)$ is linear: $u(x) = \alpha x + \beta$ for some $\alpha > 0$ and $\beta \in \boldsymbol{R}_+$

- When $u(x)$ is linear, $u(\pi x) = \pi [u(x)]$; you are indifferent between πx dollars and a probability π of x dollars (and zero dollars otherwise) – but not so when utility is strictly concave or strictly convex

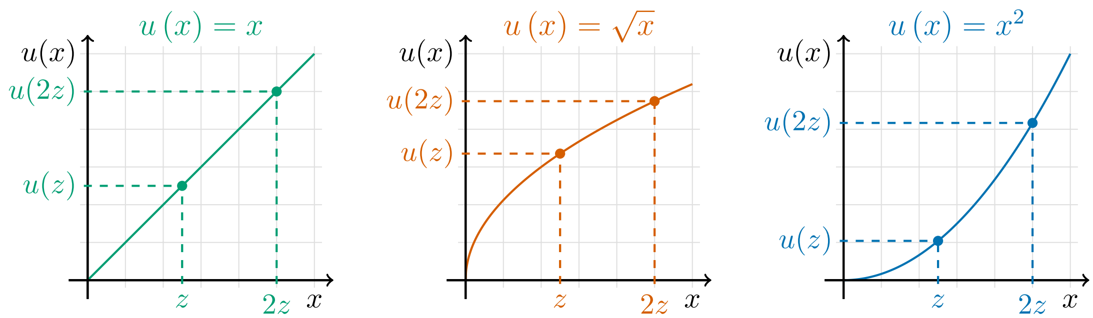

- Recall the definition of a convex function; a concave function is defined analogously

- Example above is the special case where $b = 0$, $f(b) = 0$, $a = 2z$ and $\lambda = \frac{1}{2}$

- A mathematical result called **Jensen's Inequality** is then relevant: it tells us that if $f(x)$ is a concave function, then $E[f(X)] \leq f(E(X))$. Thus, if we think of the utility that we get from a monetary payout as a concave function of the value of that payout, then most people would prefer to receive the expected value of a gamble with probability one rather than playing the lottery, and our WTP for such a risky gamble will be below its expected value.

- Define the following:

  – An individual is **risk averse** if their utility function is strictly concave: $u''(x) < 0 \Rightarrow EU(L) < u(E(L))$ for all non-degenerate lotteries

  – An individual is **risk neutral** if their utility function is linear: $u''(x) < 0 \Rightarrow EU(L) = u(E(L))$ for all non-degenerate lotteries

  – An individual is **risk loving** if their utility function is linear: $u''(x) < 0 \Rightarrow EU(L) > u(E(L))$ for all non-degenerate lotteries

- Risk neutrality: $u(x) = \alpha x + \beta$ for some $\alpha > 0$, $\beta \in \boldsymbol{R}$

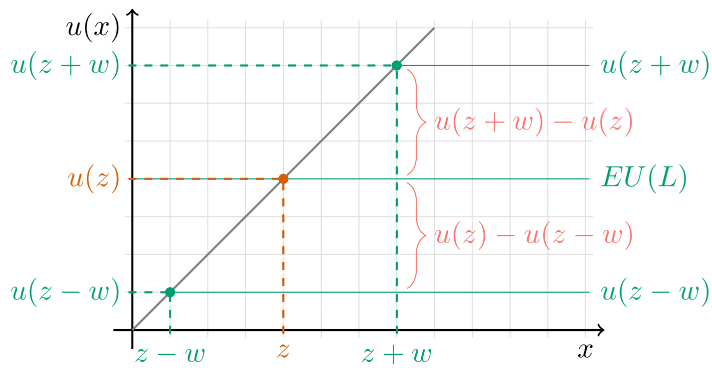

- Risk aversion:

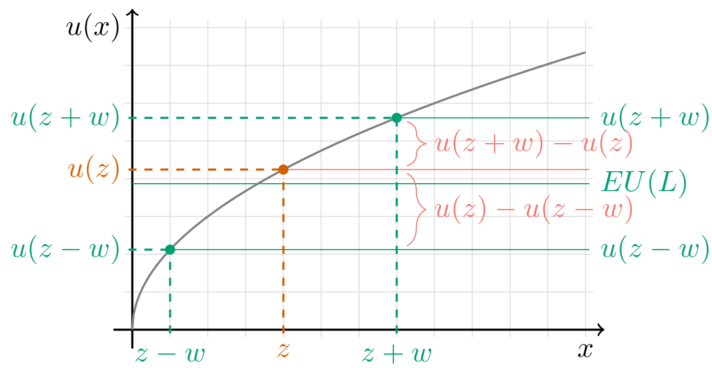

- Risk loving preferences

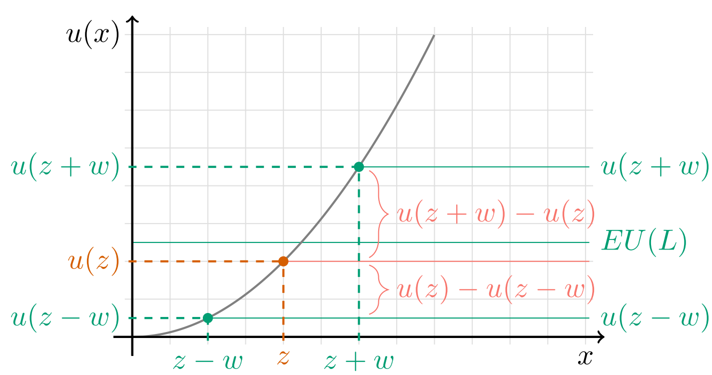

- The **certainty equivalent** of lottery L is the amount of money such that an individual is indifferent between L and $CE(L)$ dollars
  - The certainty equivalent is the maximum willingness to pay for lottery L
  - To solve for the certainty equivalent:

$$u\left(CE\left(L\right)\right) = EU\left(L\right) \Rightarrow CE\left(L\right) = u^{-1}\left(EU\left(L\right)\right) \tag{4}$$

Example: $L = \{0.5, 0.5; 0, 100\}$

Case 1: $u(x) = \sqrt{x}$ (concave utility)

$\Rightarrow EU(L) = \frac{1}{2}\sqrt{100} + \frac{1}{2}\sqrt{0} = 5 + 0 = 5$

$\Rightarrow CE(L) = u^{-1}\left(EU(L)\right) = 5^2 = \$25$

Case 2: $u(x) = x^2$ (convex utility)

$\Rightarrow u^{-1}\left(EU(L)\right) = \sqrt{\frac{1}{2}100^2 + \frac{1}{2}0^2} \approx \$70.71$

  - Notice that the certainty equivalent is *below* the expected value of the lottery for the individual with a concave utility function, but the certainty equivalent is *above* the expected value of the lottery for the individual with the convex utility function.

- The following statements are true:
  - An agent is **risk averse** $\Leftrightarrow u''(x) < 0 \Leftrightarrow CE(L) < EV(L)$ for all L
  - An agent is **risk neutral** $\Leftrightarrow u''(x) = 0 \Leftrightarrow CE(L) = EV(L)$ for all L
  - An agent is **risk loving** $\Leftrightarrow u''(x) > 0 \Leftrightarrow CE(L) > EV(L)$ for all L

8.5 Measuring Risk Aversion

- When we consider a specific lottery or risky prospect, it is reasonable to use an individual's certainty equivalent as a comparative measure of their degree of risk aversion – but how can we compare risk attitudes more generally?

- Two measures of risk attitudes:
  - **Arrow-Pratt coefficient of absolute risk aversion (CARA)**:

$$a(x) = -\frac{u''(x)}{u'(x)} \tag{5}$$

  - **Arrow-Pratt coefficient of relative risk aversion (CRRA)**:

$$r(x) = -\frac{xu''(x)}{u'(x)} \tag{6}$$

  - Numerator tells us about degree of concavity
  - Denominator normalizes units

- When the degree of absolute risk aversion does not vary with x, we say that utility is of the **constant absolute risk aversion** (CARA) form: $u(x) = -e^{-ax}$ where a is the CARA coefficient.
  - Confirm that $a(x) = a$ using the formula
  - For what values of a is an individual with a CARA utility function risk averse? Risk neutral? Risk loving?

- Intuitively, having a utility function of the CARA form means that individuals evaluation of a lottery does not depend on her level of income/wealth/consumption x. So, if she would prefer a 50/50 lottery where she is equally likely to gain five dollars or lose one dollar when her consumption is one hundred dollars, she will also prefer the lottery if her consumption drops to only a dollar. To see this, to see this, note that an individual with the utility function $u(x) = -e^{-ax}$ is indifferent between doing nothing (i.e. keeping her initial wealth $w > 0$) and playing lottery $L = \{0.5, 0.5; w - 1, w + 5\}$ when the following condition holds:

$$
\begin{aligned}
u(w) &= EU(L) \\
\Leftrightarrow -e^{-aw} &= \frac{1}{2}\left(-e^{-a(w-1)}\right) + \frac{1}{2}\left(-e^{-a(w+5)}\right) \\
\Leftrightarrow -e^{-aw} &= \frac{-1}{2}\left(e^{-aw+a}\right) + \frac{-1}{2}\left(e^{-aw-5a}\right) \\
\Leftrightarrow e^{-aw} &= \frac{1}{2}\left(e^{-aw}e^{a}\right) + \frac{1}{2}\left(e^{-aw}e^{-5a}\right) \\
\Leftrightarrow e^{-aw} &= \frac{1}{2}\left(e^{-aw}e^{a}\right) + \frac{1}{2}\left(e^{-aw}e^{-5a}\right) \\
\Leftrightarrow 1 &= \frac{1}{2}\left(e^{a}\right) + \frac{1}{2}\left(e^{-5a}\right)
\end{aligned} \tag{7}
$$

Since this expression does not depend on w, one's willingness to choose lottery L does not depend on one's level of background wealth, w. This is what is meant by constant absolute risk aversion.

- Constant absolute risk aversion may or may not be a realistic property of individual behavior. It is reasonable to assume that extremely wealth individuals will turn down positive expected value prospects to avoid the risk of a modest monetary loss? Should we expect wealthy individuals to be no more willing to take financial risks than those close to subsistence, for whom monetary losses could prove ruinous?

- If we instead believe that the degree of **relative** risk aversion does not vary with x, we say that utility is of the **constant relative risk aversion** (CRRA) form: $u(x) = \frac{x^{1-r}}{1-r}$ where r is the CRRA coefficient that indexes the degree of risk aversion.
  - Show that $r(x) = r$ when utility takes the CRRA form
  - Parallel with the CES utility function

- When $r > 0$, an individual with a CRRA utility function is risk averse; when $r < 0$, an the individual is risk loving.

- CRRA utility is characterized by **diminishing absolute risk aversion**: $a(x) = r/x$, which is decreasing in x

8.6 Criticisms of Expected Utility

8.6.1 Some Assumptions Required (for an EU Representation)

- Preferences (complete and transitive) over outcomes
- Independence: suppose L_1 is preferred over L_2 (i.e. $EU(L_1) > EU(L_2)$); then for all lotteries L_3 and for all $\alpha \in (0,1)$:

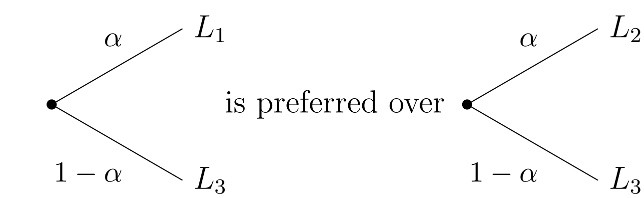

- Continuity: suppose L_A is preferred over L_B for all $\alpha \in (0,1)$:

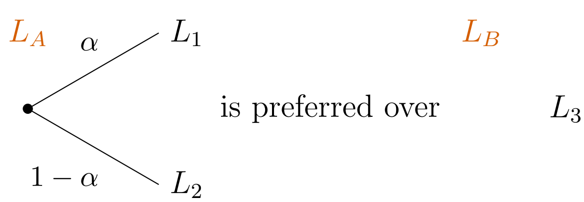

  Continuity $\Rightarrow$ if this is true, then L_1 and L_2 are both preferred over L_3

- Intuitively, continuity means that preferences over lotteries don't change suddenly at probabilities 0 and 1; independence means that when we compare two lotteries, we need only compare the outcomes that are different between them

- If these things are true, we get a beautiful mathematical model – a refinement of Bernoulli's suggestion – to use to study decisions involving risk
- Unfortunately, these assumptions are probably not (always) valid

8.6.2 Allais' Paradoxes

Which would you prefer?

- Lottery A: 1 million dollars (with certainty)
- Lottery B:
  - 5 million dollars with probability 0.1
  - 1 million dollars with probability 0.89
  - 0 dollars with probability 0.01

Which would you prefer?

- Lottery C: 1 million dollars with probability 0.11
- Lottery D: 5 million dollars with probability 0.10

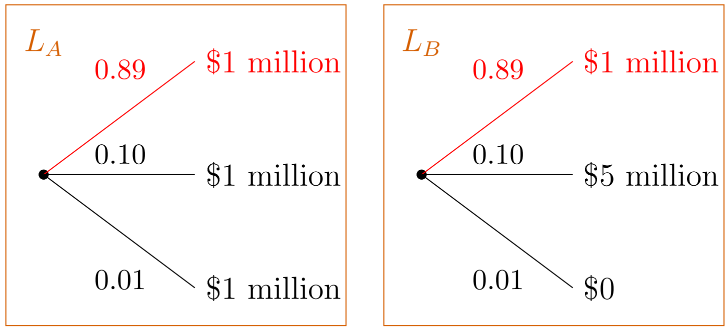

Which would you prefer?

- Lottery W: 1 million dollars with probability 1
- Lottery X: 5 million dollars with probability 0.98

Which would you prefer?

- Lottery Y: 1 million dollars with probability 0.01
- Lottery Z: 5 million dollars with probability 0.0098

$$
\begin{aligned}
EU\left(L_W\right) &> EU\left(L_X\right) \\
\Leftrightarrow u(1M) &> 0.98\left[u(5M)\right] \\
\Leftrightarrow u(1M)/100 &> 0.98\left[u(5M)\right]/100 \\
\Leftrightarrow 0.01\left[u(1M)\right] &> 0.0098\left[u(5M)\right] \\
\Leftrightarrow EU\left(L_Y\right) &> EU\left(L_Z\right)
\end{aligned}
$$

8.6.3 Alternatives to Expected Utility Theory

- ∘ Prospect Theory

- Different models for certainty vs. uncertainty