

Chapter 5: Expected Utility Theory

ECON 5001

Prof. Paul J. Healy

Based on: *Game Theory* by Giacomo Bonanno

Updated 2026-10-13 at 09:34:26

Beyond Risk Neutrality

Chapter 4: chance moves give us **lotteries** over outcomes

Initial assumption:

- All outcomes are money
- Players are **risk neutral** (rank lotteries by expected value)

We know that's not often true:

- People buy insurance. And lottery tickets!
- Many outcomes aren't money at all
 - Win/lose/draw, getting the job, where to eat dinner...

Our goal:

1. Rank lotteries over **any** kind of outcome
2. Describe different attitudes toward risk

MobLab: “Bomb Risk Elicitation Task” (BRET)

- MobLab time!
- There are 100 boxes. One of them hides a **bomb**
 - Every box is equally likely to have it
- You choose how many boxes to open (anything from 0 to 100)
- Then we find out where the bomb was:
 - **Not** in your boxes: you earn 1 point per box you opened
 - **In** your boxes: you earn 0

Feel free to talk! And play it the way you'd actually play it.

Lotteries

Lotteries

A **lottery** L is a list of outcomes and their probabilities:

$$L = \begin{pmatrix} o_1 & o_2 & \cdots & o_n \\ p_1 & p_2 & \cdots & p_n \end{pmatrix}$$

Handwritten notes:
← outcomes (pointing to o_1, o_2)
← probabilities (pointing to p_1, p_2)
 $\sum_{i=1}^n p_i = 1$
 $p_i \in [0, 1]$

The outcomes can be **anything**:

- Money: $L_1 = \begin{pmatrix} \$10 & \$20 \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}$

- Not money: $L_2 = \begin{pmatrix} \text{Buckeyes beat Michigan} & \text{Buckeyes lose to Michigan} \\ \frac{3}{5} & \frac{2}{5} \end{pmatrix}$

A sure thing is a lottery too. Just write o for “outcome o with probability 1”

Expected Value: Numbers Only

Recall the **expected value** of a money lottery:

Average earnings

$$E[L] = \underline{p_1}x_1 + \underline{p_2}x_2 + \cdots + p_nx_n$$

$$L_1 = \begin{pmatrix} \$10 & \$20 \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}$$

$$E[L_1] = \underline{\frac{1}{2}10 + \frac{1}{2}20 = 5 + 10 = 15}$$

$$L_2 = \begin{pmatrix} \text{Beat Michigan} & \text{Lose to Michigan} \\ \frac{3}{5} & \frac{2}{5} \end{pmatrix}$$

$$E[L_2] = \text{Doesn't exist}$$

Outcomes aren't numbers

Expected Value: Numbers Only

Recall the **expected value** of a money lottery:

$$E[L] = p_1x_1 + p_2x_2 + \cdots + p_nx_n$$

$$L_1 = \begin{pmatrix} \$10 & \$20 \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix} \quad E[L_1] = \underline{\hspace{2cm}}$$

$$L_2 = \begin{pmatrix} \text{Beat Michigan} & \text{Lose to Michigan} \\ \frac{3}{5} & \frac{2}{5} \end{pmatrix} \quad E[L_2] = ???$$

Expected Value: Numbers Only

Recall the **expected value** of a money lottery:

$$E[L] = p_1x_1 + p_2x_2 + \cdots + p_nx_n$$

$$L_1 = \begin{pmatrix} \$10 & \$20 \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix} \quad E[L_1] = \underline{\hspace{2cm}}$$

$$L_2 = \begin{pmatrix} \text{Beat Michigan} & \text{Lose to Michigan} \\ \frac{3}{5} & \frac{2}{5} \end{pmatrix} \quad E[L_2] = ???$$

Lesson: We can calculate an expected value when outcomes are **numbers** (like money).
When outcomes are **not** numbers, expected value doesn't exist!

So how do we rank lotteries like L_2 ?

Expected Utility

Two Kinds of Utility

Step 1: Give every *outcome* a number

- $u(o)$ is the **outcome utility** of outcome o subjective (small u : outcomes)
little u

Step 2: Use those to give every *lottery* a number

- $U(L)$ is the **expected utility** of lottery L (capital U : lotteries)
big U

Expected utility: average the outcome utility, weighting by the probabilities.

$$L = \begin{pmatrix} o_1 & o_2 & \cdots & o_n \\ p_1 & p_2 & \cdots & p_n \end{pmatrix} \mapsto \begin{pmatrix} u(o_1) & u(o_2) & \cdots & u(o_n) \\ p_1 & p_2 & \cdots & p_n \end{pmatrix}$$

big U ↓ little u ↓

$$\underbrace{U(L)}_{\text{expected utility of the lottery}} = \underbrace{p_1 u(o_1) + p_2 u(o_2) + \cdots + p_n u(o_n)}_{\text{probability-weighted average of the outcome utilities}}$$

expected value

Bonanno writes $U(o)$ for our $u(o)$, and $E[U(L)]$ for our $U(L)$.

Example

Alex's outcome utility over dinner outcomes:

Outcome o	Cazuela's	Yoshi Sushi	Dining hall
outcome utility $u(o)$	10	6	0

Two lotteries:

$$L_A = \begin{pmatrix} \text{Yoshi Sushi} \\ 1 \end{pmatrix}$$

Handwritten: 6 above Yoshi Sushi, 1 below Yoshi Sushi, and an arrow pointing to 1 with the word "lost" below it.

$$L_B = \begin{pmatrix} \text{Cazuela's} & \text{Dining hall} \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}$$

Handwritten: 10 above Cazuela's, 0 above Dining hall, and 1/2 below both.

Handwritten: Is $L_A > L_B$ or $L_B > L_A$?
 $U(L_A)$ vs. $U(L_B)$

$$U(L_A) = 1 \cdot 6 = 6$$

$$U(L_B) = \frac{1}{2}(10) + \frac{1}{2}(0) = 5$$

Alex prefers: $L_A > L_B$ because $U(L_A) > U(L_B)$

Note: expected utility exists even when expected value doesn't exist

Why Expected Utility?

When does expected utility actually describe peoples' preferences? *Axioms $\Rightarrow$*

Start with $\succeq$ over lotteries

*you have a $u(o)$
and $U(L) = \sum p_i u(o_i)$*

Theorem (von Neumann & Morgenstern, 1944)

If your preference $\succeq$ over lotteries satisfies four "rationality" axioms (see the book), then you have some outcome utility $u(o)$ such that $\succeq$ is represented by the expected utility function

$$U(L) = \sum_i p_i u(o_i)$$

In other words: $L \succ L' \Leftrightarrow U(L) > U(L') \Leftrightarrow \underbrace{\sum_i p_i u(o_i)}_{\text{E.U. of } L} > \underbrace{\sum_i p'_i u(o'_i)}_{\text{E.U. of } L'}$

Note: "as if" theory $\rightarrow$ *not actually calculating E.U.*

*As long as $\succeq$ satisfies
the axioms, you
behave AS IF
you're calculating
E.U.*

If you're interested: Axioms are in Bonanno §5.3 (but **not** on the exam).

From here on: we **assume** players maximize expected utility

Why Expected Utility?

When does expected utility actually describe peoples' preferences?

Theorem (von Neumann & Morgenstern, 1944)

If your preference $\succsim$ over lotteries satisfies four “rationality” axioms (see the book), then you have some outcome utility $u(o)$ such that $\succsim$ is represented by the expected utility function

$$U(L) = \sum_i p_i u(o_i)$$

In other words: $L \succ L' \Leftrightarrow U(L) > U(L') \Leftrightarrow \underbrace{\sum_i p_i u(o_i)}_{\text{E.U. of } L} > \underbrace{\sum_i p'_i u(o'_i)}_{\text{E.U. of } L'}$

Note: “as if” theory

If you're interested: Axioms are in Bonanno §5.3 (but **not** on the exam).

From here on: we **assume** players maximize expected utility

Non-EU theories

Prospect theory
weight probabilities

$$U(L) = \sum_{i=1}^n w(p_i) \cdot u(o_i)$$

↑
distort probabilities

ex: overweight low
probabilities

Not on exam



John von Neumann (1903–1957)



Oskar Morgenstern (1902–1977)

- **von Neumann:** Budapest math prodigy. Princeton IAS (with Einstein)
 - Manhattan Project, modern computer design, minimax theorem
- **Morgenstern:** Vienna economist. Left Austria in 1938
- **1944:** *Theory of Games and Economic Behavior*.
- Expected utility axioms: an appendix to the 1947 edition!

Rescaling the outcome utility

Can We Rescale Any Way We Want?

Chapter 2: You can rescale utility as long as it respects the ordering $\succsim$

$$L_A = \begin{pmatrix} \text{Yoshi Sushi} \\ 1 \end{pmatrix} \quad L_B = \begin{pmatrix} \text{Cazuela's Dining hall} \\ \frac{1}{2} \quad \frac{1}{2} \end{pmatrix}$$

What if we squared the outcome utilities? Preserves the ordering...

NOT the
same $\succsim$
over
lotteries

Outcome o	Cazuela's	Yoshi Sushi	Dining hall
$u(o)$	10	6	0
$u(o)^2$	100	36	0

Original u : $U(L_A) = 1 \cdot 6 = 6$

Squared u : $U(L_A) = 1 \cdot 36 = 36$

$U(L_B) = \frac{1}{2} \cdot 10 + \frac{1}{2} \cdot 0 = 5$

$U(L_B) = \frac{1}{2} \cdot 100 + \frac{1}{2} \cdot 0 = 50$

$L_A \succ L_B$

$L_B \succ L_A$

Can We Rescale Any Way We Want?

Chapter 2: You can rescale utility as long as it respects the ordering $\succsim$

$$L_A = \begin{pmatrix} \text{Yoshi Sushi} \\ 1 \end{pmatrix} \qquad L_B = \begin{pmatrix} \text{Cazuela's Dining hall} \\ \frac{1}{2} \qquad \frac{1}{2} \end{pmatrix}$$

What if we squared the outcome utilities? Preserves the ordering...

Outcome o	Cazuela's	Yoshi Sushi	Dining hall
$u(o)$	10	6	0
$u(o)^2$	100	36	0

Original u : $U(L_A) =$

$U(L_B) =$

Squared u : $U(L_A) =$

$U(L_B) =$

When using expected utility for lotteries,
you **cannot** rescale outcome utilities arbitrarily!

Linear Rescaling

Which rescalings *are* allowed? Only **linear** ones:

$$v(o) = a \cdot u(o) + b \quad (\text{with } a > 0)$$

$\uparrow$ original $u(o)$

linear rescaling
 $\Rightarrow$
same $\geq$

- You can **multiply** by a positive number $a > 0$
- You can **add** (or subtract) a constant b

} linear transformations

Like temperature: $^{\circ}\text{F} = \frac{9}{5} \times ^{\circ}\text{C} + 32$.

squaring $u(o)$ was
NOT linear $\Rightarrow$
different $\geq$

Technically, a “positive affine transformation” (Bonanno)

Normalizing

Since linear rescaling is allowed, we can always pick a convenient scale:

$$u(\text{best outcome}) = 1$$

$$u(\text{worst outcome}) = 0$$

This is called the **normalized** outcome utility. How to get it:

1. Subtract $u(\text{worst})$ from every number
2. Divide every number by the new $u(\text{best})$

(now the worst is 0)

(now the best is 1)

Outcome	o_1	o_2	o_3	o_4
$u(o)$	8	2	-2	0
After step 1	10	4	0	2
After step 2	1	0.4	0	0.2

subtract -2 $-(-2) \rightarrow +2$
divide by 10 $(\times \frac{1}{10})$

Equivalently, scale to $[0, 100]$

Attitudes Toward Risk

Risk Averse, Risk Neutral, Risk Loving

Back to **money** lotteries, where $E[L]$ exists. Offer someone a choice:

the lottery L vs. getting $\$E[L]$ for sure

Ex: $L_1 = \begin{pmatrix} \$10 & \$20 \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}$ vs. $\$E[L_1] = \$15$ for sure. Which would you take?

risky *no risk*

Definition

- **Risk averse**: strictly prefers the sure payment of $\$E[L_1]$ (expected value for sure)
 - **Risk neutral**: indifferent between them $\leftarrow$ Expected utility = expected value
 - **Risk loving**: strictly prefers the risky lottery L_1
- $u(10) = 10$
 $u(20) = 20$

Can depend on the lotteries you face

Important: This only makes sense when outcomes are numbers (like money).

With non-monetary lotteries, there is no such thing as “risk averse” or “risk loving”

Risk Attitudes and the Shape of $u(o)$

With expected utility, the definition becomes a comparison of two numbers:



$\underbrace{U(L)}$
expected utility of the lottery

vs.

$\underbrace{u(E[L])}$
utility of the expected value

Risk averse $\iff U(L) < u(E[L]) \iff u(o)$ is **concave**

Risk neutral $\iff U(L) = u(E[L]) \iff u(o)$ is **linear**

Risk loving $\iff U(L) > u(E[L]) \iff u(o)$ is **convex**




next slides show why

Not in Bonanno, but important.


Risk Averse $\iff$ Concave

$$L = \begin{pmatrix} 10 & 20 \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}$$




L_1 pays \$10 or \$20 with equal chance. $E[L_1] = \$15$.

Risk Averse $\iff$ Concave



Expected utility $U(L_1)$ is halfway between $u(10)$ and $u(20)$: the midpoint of the line.

Risk Averse $\iff$ Concave




The curve is *above* the line, so $u(15) > U(L_1)$. In terms of preferences: $\$15 \succ L_1$

$u(E(L_1)) > U(L_1)$ risk averse

All Three Cases


Concave $u'' < 0$



$$U(L_1) < u(15)$$

Risk averse


Linear $u'' = 0$



$$U(L_1) = u(15)$$

Risk neutral

Convex $u'' > 0$



$$U(L_1) > u(15)$$

Risk loving

Red dot: expected utility of the lottery. **Blue dot:** utility of \$15 for sure.

Risk Attitudes in the Real World

Lots of real decisions boil down to “lottery vs. sure thing”:

- **Insurance.** Pay a premium for sure, or risk losing your car or house?
- **Extended warranties.** Pay a premium, or risk losing your laptop
- **Salary vs. commission.** Steady paycheck, or pay that depends on sales?
- **Index funds vs. single stocks.** Spread the risk, or bet on one company?
- **Sports betting, casinos, Ohio Lottery.** Pay to *take on* risk! Risk loving

Real behavior:

- The same person might have car insurance *and* buy lottery tickets
- People treat gains and losses differently (“Prospect Theory”)

Regardless, **the shape of $u(o)$ tells you how they handle risk**

Example: Car Insurance

Should You Insure Your Car?

Ruth and Logan have:

- \$1,000 in the bank, and a car worth \$15,000 (total wealth: \$16,000)
- A 25% chance of totaling the car this year

Nationwide offers **full insurance** for a \$6,000 premium:

- If the car is totaled, they reimburse you the full \$15,000

Final wealth (in \$1,000s):

	No crash (75%)	Crash (25%)
L_1 No insurance	16	1
L_2 Insurance	10	10

$$E\{L\} = \frac{3}{4} \cdot 16 + \frac{1}{4} \cdot 1$$

$$12 + 0.25 = 12.25$$

risky
safe $E\{L\} = 10$

$$E[\text{No insurance}] = \underline{12.25}$$

$$E\{I_{ns}\} = 10$$

A risk-neutral person would choose: No Ins.

Two Different Utility Indices

Let o be final wealth (in \$1,000s).

Ruth: $u_R(o) = \sqrt{o}$

$$U_R(\text{No ins.}) = \frac{3}{4}\sqrt{16} + \frac{1}{4}\sqrt{1} = 3.25$$

✓

$$U_R(\text{Ins.}) = u_R(10) = \sqrt{10} \approx 3.16$$

Ruth chooses: No Ins.

Logan: $u_L(o) = \ln(o)$

$$U_L(\text{No ins.}) = \frac{3}{4}\ln(16) + \frac{1}{4}\ln(1) \approx 2.0794$$

$$U_L(\text{Ins.}) = u_L(10) = \ln(10) \approx 2.30$$

Logan chooses: Ins.

Both concave $\Rightarrow$ Both risk averse

Logan is more risk averse than Ruth

Two Different Utility Indices

Let o be final wealth (in \$1,000s).

Ruth: $u_R(o) = \sqrt{o}$

$$U_R(\text{No ins.}) = \frac{3}{4}\sqrt{16} + \frac{1}{4}\sqrt{1} = 3.25$$

$$U_R(\text{Ins.}) = u_R(10) = \sqrt{10} \approx 3.16$$

Ruth chooses: _____

Logan: $u_L(o) = \ln(o)$


$$U_L(\text{No ins.}) = \frac{3}{4}\ln(16) + \frac{1}{4}\ln(1) \approx 2.0794$$

$$U_L(\text{Ins.}) = u_L(10) = \ln(10) \approx 2.30$$

Logan chooses: _____


Both are concave, so **both** are risk averse. But only one buys insurance!

Seeing It in a Graph



Both rescaled so that $u(1) = 0$ and $u(16) = 1$: $(\sqrt{o} - 1)/3$ and $\ln(o)/\ln(16)$.
Logan's utility is more concave: Logan is **more** risk averse than Ruth.


Seeing It in a Graph



No insurance: 75% chance of the best, 25% chance of the worst.

So $U(\text{No ins.}) = 0.75$ for **both** of them: the dashed line. (Before rescaling: 3.25 and 2.08)

Seeing It in a Graph



Insurance gives 10 for sure. Is $u(10)$ above or below the dashed line?

Ruth: below, so no insurance. **Logan:** above, so buy insurance.

How Much Would They Pay?

The expected loss is $\frac{1}{4} \times \$15,000 = \$3,750$. Nationwide charges \$6,000.

What's the **most** each would pay for full insurance?

- **Ruth:** find the sure wealth w with $\sqrt{w} = 3.25$
- **Logan:** find the sure wealth w with $\ln(w) = \frac{3}{4} \ln(16)$

On the graph: where the dashed line crosses each curve.

Lesson: Knowing someone is risk averse doesn't tell you what they'll choose. You need to know their outcome utility.

Let's Work One Out

The Prize Wheel

Carmen gets one spin of a prize wheel at the Ohio Union. The possible prizes:

- M : a ticket to the Michigan game J : a pint of Jeni's ice cream
- D : a dozen Buckeye Donuts N : nothing

Carmen's ranking: $M \succ J \succ D \succ N$. Carmen also tells us:

- J for sure is just as good as a 50% chance of M (otherwise N)
- D for sure is just as good as a 20% chance of M (otherwise N)

Carmen can pick which wheel to spin:

$$W_1 = \begin{pmatrix} J & D \\ 0.5 & 0.5 \end{pmatrix} \quad W_2 = \begin{pmatrix} M & N \\ 0.3 & 0.7 \end{pmatrix} \quad W_3 = \begin{pmatrix} M & J & D \\ 0.1 & 0.4 & 0.5 \end{pmatrix}$$

The Prize Wheel

1. Write down Carmen's normalized outcome utility.
2. Find the expected utility of each wheel. Which does Carmen spin?
3. Rescale using $v(o) = 10u(o) + 5$. Does the ranking of wheels change?
4. Now rescale using $w(o) = u(o)^2$. Does the ranking change? Why?
5. Is Carmen risk averse, risk neutral, or risk loving?

What Have We Learned, Charlie Brown?

- **Expected value** needs numeric outcomes. **Expected utility** doesn't
- $u(o)$: **outcome utility** over outcomes
- $U(L)$: **expected utility** over lotteries
 - $U(L) = p_1u(o_1) + \dots + p_nu(o_n)$
- The outcome utility can only be rescaled **linearly**
 - So we can normalize: best outcome = 1, worst outcome = 0
- Risk averse / neutral / loving $\iff u(o)$ is concave / linear / convex
 - Only for numeric outcomes!
- Two risk-averse people can still make different choices

Next: games where the payoffs are expected utilities

- Bonanno, *Game Theory*, Chapter 5
 - §5.1 (risk attitudes) and §5.2 (expected utility)
 - §5.3 (the axioms): optional, read it for fun
- Remember: Bonanno's $U(o)$ is our $u(o)$. His $E[U(L)]$ is our $U(L)$
- Exercises: §5.4.1 (risk attitudes) and §5.4.2 (expected utility)
 - Exercises 5.2 and 5.4: outcomes that aren't numbers
 - Exercise 5.11: two risk-averse people who choose differently
- Solutions in §5.5